AFGL - TR - 81 -0304

LEVEL II

A MEASUREMENT OF LONG TERM TILT IN
COLORADO AND WYOMING

J. Christopher Harrison
Judah Levine

University of Colorado
Campus Box 440
Boulder, Colorado 80309

12 49

AD A108865

1 September 1981

Final Report for the period 6 December 1977 - 31 May 1981

Approved for public release; distribution unlimited

DTIC FILE COPY

AIR FORCE GEOPHYSICS LABORATORY
AIR FOR SYSTEMS COMMAND
UNITED STATES AIR FORCE
HANSCOM AFB, MASSACHUSETTS 01731

088400

8 112 28 067

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: AFGL-TR-81-0304	2. GOVT ACCESSION NO.: AD-A108865
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): A Measurement of Long Term Tilt in Colorado and Wyoming	5. TYPE OF REPORT & PERIOD COVERED: Final Report (6 Dec 77-31 May 81)
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): J. Christopher Harrison, Judah Levine	8. CONTRACT OR GRANT NUMBER(s): F 19628-78-C-0065
9. PERFORMING ORGANIZATION NAME AND ADDRESS: University of Colorado, Campus Box 440, Boulder, Colorado 80309	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 61102F, 2309G2AB
11. CONTROLLING OFFICE NAME AND ADDRESS: Air Force Geophysics Laboratory, Hanscom Air Force Base, Massachusetts 01731, Monitor: G. H. Cabaniss/LWH	12. REPORT DATE: 1 September 1981
	13. NUMBER OF PAGES: 48
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Borehole tiltmeters
Earth tides
Secular tilt
Tidal analysis
Tilt coherence

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

We have developed a borehole tiltmeter capable of operating for long periods of time at tidal sensitivity, and we have deployed six of these tiltmeters at two sites near Boulder, Colorado. The data from the tiltmeters are recorded using a data acquisition and control system capable of recording the data from a large number of widely spaced instruments linked to a central site by means of telephone lines or radio links. We have also written analysis software to

Unclassified
SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Abstract, Block 20, Continued
examine the power spectra of the tilt records, to extract the tidal admittance from time series and to evaluate the coherence between different instruments as a function of frequency. We have used our tilt array to measure the amplitude and phases of the major components of the earth tides, and we have examined the time dependence of the tidal admittance. We also have some preliminary estimates of the coherence between independent instruments at periods longer than one day.

Unclassified
SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

SUMMARY OF OBJECTIVES

1. To conduct field and laboratory measurements of long term crustal motions at periods longer than four minutes.

2. To install two borehole instrument systems about 1 km apart near Boulder, to investigate their coherence and evaluate the influence of local effects.

3. To develop and deploy an array of instruments in Wyoming, Montana or Colorado.

INTRODUCTION

The measurement of crustal tilt can provide information on the elastic constants of the earth and on the response of the earth to stresses of tectonic origin. It can also be used to study strains in areas of seismic activity, and might conceivably provide insight into precursory phenomena.

In order to make unambiguous tilt measurements, it is essential that the tilt sensor be both sensitive and stable, and that it be insensitive to environmental perturbations. Furthermore, the instruments must be installed in such a way that the reported tilts are as uncontaminated as possible by various local effects produced by rainfall or by other extraneous disturbances.

The success of a given design can be most easily evaluated by comparing the tilt records obtained from several closely spaced instruments, since it is unreasonable to expect meaningful records from instruments which show only weak coherence when operated a short distance apart.

An equally stringent test is a comparison of tidal tilt data with theory. Although the exact amplitude and phase of the admittance may not be known at a particular site, a tidal analysis can be used to study the stability of the instrumental gain, the linearity of the system, and the sensitivity of the system to the diurnal temperature cycle.

The following is a list of the names of all of the people who have worked on the project. Except for the first three, all of the people listed have worked on this project on a temporary or part-time basis.

J. Christopher Harrison, Prof. of Geology, Co-P.I.
Judah Levine, Adj. Prof. of Physics, Co-P.I.
Charles Meertens, Research Assistant (Graduate Student)
Charles G. Hansen, Undergraduate Research Aide
John Van Zant Harvey, Research Assistant (Graduate Student)
Robin M. Jeffries, Computer Programmer
John Magyar, Postdoctoral Research Associate
Mark D. Pearson, Undergraduate Research Aide
Michael G. Schnapp, Research Assistant (Graduate Student)
Jeanette A. Trebing, Secretary
Lynn R. Walloch, Secretary

BOREHOLE DESIGN

The basic design of the borehole has not changed much since our first holes described in our annual report. The holes we use are nominally 15 cm in diameter and 33 m deep. After the hole is drilled, a steel casing, 135 mm in diameter with 6 mm walls, is pressed into the hole. The casing is shipped to the site in 6-meter lengths, and is welded into a continuous, water-tight pipe as it is lowered down the hole. The casing is sealed in place by means of cement poured down to the bottom of the hole before the casing is inserted and around the sides of the casing after it is in place.

The steel casing terminates at the bottom in an 2.3-meter long stainless-steel section used to hold the tiltmeter capsule. The bottom section is 11.5 cm in outside diameter, and has walls 6 mm thick. The top of this section has a transition section to the standard carbon steel pipe which makes up the rest of the casing.

The bottom section is closed at its lower end by a plate welded on in the shop. A hemispherical knob is welded to the inside of the bottom plate to support the weight of the tiltmeter capsule.

INSTRUMENT CAPSULE

The instrument capsule is a 1.8-meter length of stainless steel tubing closed at the bottom and having a pair of contact points and a flat spring welded on near its top and a second pair with a second flat spring near its bottom. The top of the capsule is sealed with a cap attached by screws and containing an O-ring. The cap has a water-tight opening for the electrical cable to pass through, hooks for attaching the lifting cable, and a post for attaching the orienting rods (see below).

The capsule was designed to minimize tilt-strain coupling due to cavity effects. Harrison (1976) has shown that there is no cavity effect if the side of the borehole is used as a reference axis for the tiltmeter, and only a small effect if the center of the bottom of the hole is used as one reference point.

The capsule is raised and lowered by hand using a stainless-steel lifting cable attached to the top cap. The capsule can be raised or lowered in a few minutes.

The length of the capsule also improves the coupling between the tiltmeter and the earth by increasing the effective lever arm from a few centimeters (the length of the tiltmeter) to approximately 2 meters (the length of the capsule).

In fig. 1 we show the capsule installed at the bottom of the borehole.

The capsule weighs approximately 30 lbs. This weight is supported on the hemispherical knob located at the bottom of the casing. The cables connecting the capsule with the surface are left slack; their weight is supported from a bracket at the top of the casing.

CAPSULE ORIENTATION

In order to compare the tilts recorded by different instruments or to compare the tidal tilt with theory, it is necessary to know the orientation of the tiltmeter. Magnetic sensors inside the capsule to sense the direction of the earth's magnetic field cannot be used since the casing is magnetic. The capsule itself is usually not visible from the surface so that we cannot determine the orientation by sighting on the capsule from the top of the hole.

We have developed a system involving a series of light rods to determine the orientation of the capsule at the bottom of the hole. A post is welded to the top cap of the tiltmeter capsule with a flat side oriented so that the normal to the flat side is along the axis of one of the tiltmeters. As the capsule is lowered, additional sections of rod are added. The first section attaches to the post on the capsule using a trapped ball and a detent to provide an easily removeable coupling; subsequent sections bolt together with small screws. Each section is notched so that it can only be attached in one orientation. After the orientation of the top notch is determined using a transit and a compass, the entire series of rods is removed from the capsule by simply lifting gently so as to disconnect the bottom rod from the tiltmeter capsule.

This method has been used to determine the azimuth of instruments at the bottom of holes 33 m deep. We estimate that the uncertainty in our determination of the azimuth is about one degree, although we have no independent way of checking it. It is difficult to extend this technique to holes much deeper than 33 meters since the weight of the rods becomes appreciable, and the capsule becomes very difficult to handle.

LEVELLING PLATFORM

The levelling platform mounts inside of the tiltmeter capsule and is used to support and level the tilt sensors. It is supported on a three-point kinematic mount. Two of the supports are screws connected to small motors. These motors are driven from the surface and are used to re-zero the instrument. The levelling platform has a range of approximately five degrees in any direction, so that the instrument must be vertical to within five degrees to start with. This limited re-zeroing capability imposes some constraint on the drillers, but the required tolerance can be met in 33 meter deep holes.

In fig. 2 we show the tiltmeters and the levelling platform.

TILT SENSORS

All of our instruments use mechanical tilt sensors. The sensors are pendulums mounted on fine wire suspensions. All of them have been manufactured by Larry Burris of Instech, Inc. They have natural periods of a fraction of a second (the actual periods differ slightly since both horizontal and vertical pendulums have been used). In all cases the pendulum is placed between two fixed plates separated by about one mm.

The two capacitances (from the center plate to either end plate) form two arms of a capacitance bridge. The other two arms are formed by two fixed impedances, traditionally the center-tapped secondary of a carefully balanced transformer. We have used two precision resistors to complete the bridge rather than the two halves of the transformer since this configuration is easier to adjust

If the two outer plates are driven out-of-phase by a high frequency a.c. signal (of order 10 kHz), the signal at the center plate has a magnitude proportional to the deflection of the center plate from the electrical midpoint of the system (i.e. from that position at which the potential in the gap is the same as the mean of the outer plate driving voltages), and a phase (relative to the driving signal) giving the direction of the deviation.

TILT ELECTRONICS

Since both the amplitude and phase of the signal at the center plate carry information we use a phase-sensitive detector to derive a voltage proportional to the tilt. The reference for the phase detector is obtained from the drive signal to the outer plates. The signal input to the phase-sensitive detector is driven by a preamplifier having a gain of several thousand (the actual gain is set by running the tilt sensor on a test table and adjusting the gain to yield an overall sensitivity at the output of nominally 2 volts/microradian).

The low mechanical sensitivity of the pendulums requires a correspondingly high gain in the electronics if the instrument is to be sensitive to tilts at the nano-radian level.

An equally difficult problem is the sensitivity of the system to fluctuations in the parasitic capacitances in the circuit (especially in the components that drive the outer plates) and between the connecting wires and ground. The capacitance between the center plate and either end plate is of order 3 pf, and a full-scale tilt (of order 5 micro-radians) results in a change of only about one part per million in this capacitance. Thus the entire system must have stray capacitances whose values change by less than a small fraction of a pf, and in practice this requirement is far more stringent than the need for high electrical gain. We have addressed this requirement by keeping the system as small as possible and by placing the entire electronics package in the temperature-stabilized environment at the bottom of the hole.

The most sensitive part of the circuit is the part that generates the drive voltage for the outer plates. The drive voltage for the outer plates is obtained from a small step-up transformer. In many capacitance bridge circuits, the transformer is center-tapped with the center-tap grounded. The bridge is formed by the two capacitances of the tilt sensor and the two halves of the transformer secondary. Thus the tilt signal is obtained by comparing the two capacitances against the two halves of the transformer secondary. The stray capacitances between turns are appreciable, and it is difficult to find commercial transformers for which the manufacturer will guarantee the stability of the secondary impedance. In addition the resistance of each half of the secondary is not negligible, and this resistance combines with the inductance of the coil and the stray capacitance to ground to form a complex frequency-dependent driving impedance for the outer plates. These impedances do not cause problems if they are carefully balanced. However it is difficult to match the temperature coefficients very well. Our circuit therefore uses a secondary with no center-tap driving a pair of matched resistors to ground. In this circuit the bridge is formed between the two capacitances of the tilt sensor and the two precision resistors. The driving impedance of the transformer is unimportant, provided only that it is small enough to act like a pure voltage source when driving the bridge. In this circuit, the balance may be trimmed by inserting a small potentiometer if necessary. Although we have not found this to be necessary, it has been necessary to adjust the relative phase of the two outputs so that the two plates are driven exactly out of phase. (This lack of perfect anti-symmetry presumably arises from a slight difference in the stray capacitance between each half of the secondary and ground.) Unless this is done quite carefully, the signal at the input stage does not go to zero even at the electrical center due to the appreciable quadrature voltage. The quadrature voltage may be large enough to saturate the input stage (if the system is sufficiently out of balance) thereby limiting the sensitivity of the system. The trimming capacitors required to cancel the quadrature voltage are on the order of a few pf, so that lead dress is critical. The size of the trimming capacitance required is determined by trial and error for each instrument after it has been assembled. The optimum trimming capacitance is that value which minimizes the quadrature voltage at null.

Although the quadrature component of the imbalance is serious since it limits the a.c. gain of the system, it does not contribute to the final tilt output. The in-phase component is much more important in that it adds to the true tilt signal and produces an offset whose stability must be controlled if the full stability of the system is to be realized. The stability of this offset is difficult to assess since it depends on parasitic capacitances that are difficult to measure.

The entire circuit is carefully enclosed in a shielded box to minimize changes in the stray capacitances. Nevertheless it is possible that some of our long-term 'tilts' and some of the residual sensitivity to temperature changes (both of these effects will be discussed later) may in fact be caused by changes in the values of some of these components with a resulting change in the balance point of the capacitance bridge.

The tiltmeter electronics package consumes approximately 3 watts of electrical power. This power is dissipated as heat within the capsule. In our early installations, the main heat loss mechanism was via air convection within the capsule. As the air near the electronics was warmed, an inversion was established as cooler air above trapped this warm air. After about twenty minutes, the warm air would become hot enough to displace the cool air, the temperature at the electronics would drop and the process would be repeated. This periodic convection process produced a corresponding periodic "tilt." We solved this problem in our later designs by surrounding the electronics with baffles to prevent this convection process and to increase the rate of cooling via conduction to the capsule walls. These baffles are made out of foam rubber; they are cut to fit snugly around the electronics.

POWER SUPPLY

The electronics package described above together with the digitizing system to be described below consume approximately 7 watts of power. This is somewhat more than can be conveniently supplied by batteries. All of our existing sites have commercial power nearby, and this will be a restriction on site selection for the forseeable future.

In order to minimize coupling between the tiltmeter and surface electrical noise, we do not have the tiltmeter power supply on the surface. The power supply on the surface produces +5 v. d.c. only, and this is converted inside the hole to the +/- 15 volts required for the tiltmeter. In addition to decoupling the tiltmeter from surface noise it also allows us to define a totally independent ground system for the tiltmeter and for the surface equipment.

The voltage drop in the cable between the surface power supply and the power converter is removed by using remote-sensing supplies with a separate set of leads to sense the voltage at the electronics. The 5 volts itself is used only for digital TTL circuitry, and noise pickup is not a problem.

Almost all of the low voltage wiring is inside of the borehole casing which is a heavy, low-inductance path to ground. This wiring is probably immune from lightning strikes. However the commercial power input and the transmission cables are exposed. They are both protected with MOV type surge protectors, but these are unlikely to be of much use in disipating the energy from a nearby lightning strike. Both of our sites are near large antenna arrays, and these probably protect us from direct hits. In fact we have had more trouble with lightning strikes on the telephone lines (which have destroyed modems and telephone company junction boxes) and on the JILA building itself. These strikes have caused problems with the central acquisition computer.

FEEDBACK SYSTEM

The electronics package described above produces a voltage proportional to tilt. The actual calibration depends on the gain of the electronics and on the mechanical sensitivity of the system. While the latter is determined by the physical size of the pendulum and is probably stable in time, the former may change as the components age.

It is possible to overcome this change in calibration by using feedback. The output of the electronics package is fed back to the tilt sensor to drive the plate back to the electrical center of its travel. Since the electrostatic force between two capacitor plates is always attractive, the feedback voltage must be superimposed on a bias so that the plate can be driven either way. Some systems apply the bias by raising the pendulum above d.c. ground. This has the disadvantage that the d.c. bias voltage is applied directly to the input of the pre-amplifier which, if it is constructed as a classical current-amplifying stage, is sensitive to d.c. We have chosen instead to apply the d.c. bias voltage to the outer plates, and to sum this bias voltage with the tilt error signal using conventional operational amplifiers. The sum of the error voltage and the bias voltage is applied to one outer plate while the difference is applied to the other. In this way the plate can be driven back towards the electrical center of the system.

As with all servo loops, the phase of the feedback must be tightly controlled. Since the pendulum will show an appreciable phase shift in its response as its resonant frequency is approached, it is important to confine the feedback to frequencies well removed from the natural period of the pendulum. In our case this means that the feedback system cannot be used at periods shorter than tens of seconds, since otherwise the phase shift of the response (relative to its low frequency value) will tend to drive the system into a sustained oscillation at a frequency near the resonant frequency of the pendulum. The limiting of the feedback frequency response is accomplished with a simple lowpass filter with a corner frequency at six seconds. Since this period is much longer than the resonant period of the pendulum, this lowpass filter dominates the system response, and the feedback on the pendulum is confined to the linear, zero-phase-shift domain. High frequency disturbances are therefore attenuated by the lowpass filter and are not fedback to the pendulum. At these high frequencies the pendulum is not servoed and behaves like a classical, free system. In order for this system to work, the noise at these non-fedback frequencies must be small enough so that the pendulum does not move far from its equilibrium position or else the input stage will saturate on the high frequency noise peaks, thus essentially shutting down the loop. Feedback is therefore only effective at quiet sites with not much noise above about 0.5 Hz.

It is important to understand the advantages and disadvantages of a fedback system, since feedback is often promoted as a universal cure. The primary advantage of feedback is that the gain of the electronics becomes unimportant, provided only that it is above a nominal minimum. The force necessary to drive the plate back to the electrical center of its travel depends on the physical properties of the pendulum, and the relationship between the output signal (the correction voltage) and the tilt is therefore independent of the details of the amplifer. A second advantage arises from the fact that the pendulum is always kept at the same point relative to the outer plates so that non-linearities in the response produced by capacitance fringe fields or by friction in the support are minimized. (This second advantage is of somewhat less importance in our system since the pendulum hardly moves in any case. A full scale deflection of 5 micro-radians corresponds to a motion of the pendulum of only 250 nm.

a bit less than the wavelength of blue light.)

However, these advantages are not achieved at no cost. The feedback system always drives the pendulum back to the electrical center of its travel, i.e. to that point at which the in-phase signal at the center plate is zero. Thus while the feedback concept attenuates changes in gain, it does not attenuate changes in d.c. offset, i.e. in the determination of when electrical zero is reached.

In a conventional (i.e., non-fedback) system, the electronics process the a.c. signal from the center plate and produce a d.c. voltage proportional to its magnitude and phase. Most of the circuit operates at the a.c. drive frequency, and the circuit is therefore insensitive to small d.c. offsets. The gain is set by the ratios of various resistors, rather than by less reliable active elements. In a well designed conventional system offsets are not much of a problem, since the high gain stages are not sensitive to d.c. and the d.c. sensitive stages are at the end of the chain (where their offsets will not be multiplied by a following stage) and are of relatively low gain.

This is not true in a fedback system, however. The bias voltage and the error voltage must be processed as d.c. signals, and drifts in either the bias supply or in the summing amplifiers appear as first order changes in the zero and are not taken care of by the feedback principle. Thus in return for removing the sensitivity to gain changes in the main amplifier the feedback system imposes extremely stiff requirements on the d.c. part of the loop. In addition, changes in the magnitude of the bias voltage produce changes in the gain of the system.

Changes in offset or changes in d.c. bias voltage are produced by changes in the active elements in the summing circuit (i.e. in the op-amps used for the summer and in the zener diodes or batteries used for the bias voltage), and are not easily controlled. All of these effects have significant temperature coefficients, so that the loop must be operated in a controlled environment.

In summary, feedback will improve the linearity of an instrument, but its effects on the long term stability and on the calibration are not clear.

We have been experimenting with one feeback system to gain some experience with feedback principles. As we will show below, we see no non-linearities in our open-loop sensors, and we would therefore expect the fedback instrument to be not much better than the open-loop devices. Although the complete story is not yet in, this in fact seems to be the case.

DIGITIZING SYSTEM

Every hole has its own digitizer. This has several advantages. The most important advantage is that it eliminates the need for transmitting analog tilt information on the surface where it is subject to noise pickup. It also increases the isolation between holes and prevents large ground loops between widely spaced instruments.

The digitizer in each hole is relatively simple. It is made up of an analog multiplexor, an analog-to-digital converter, and a control section made up of approximately 25 integrated circuits.

The digitizer communicates with the recording system via two optically isolated coaxial cables using a bit-serial transmision system. (It would also be possible to use twisted pair cables, but the decrease in cost of the cable would be offset by the increase in the complexity of the cable drivers and receivers. Coaxial cable can be driven from ordinary open-collector power drivers, but twisted pair requires true differential transmitters and receivers to be fully effective.) The transmission system is the same as ordinary RS-232C except that the voltage levels are standard TTL levels rather than the bi-polar levels used in the RS-232C protocol. By optically isolating the cables, we decouple the instrument from the ground at the recording system which may be several hundred meters away.

The digitizer remains idle until it is activated by a command from the recording system. Commands take the form of 8-bit characters transmitted at 300 baud on one of the coaxial cables. For convenience, standard ASCII characters are used, although in fact any 8-bit codes would do. The control section recognizes three commands: Reset multiplexor to channel zero (ASCII code %); Advance multiplexor to next channel (ASCII code *); begin conversion and report value when completed (ASCII code ?). The reset and advance commands are used to set the multiplexor to any one of eight analog inputs. After the multiplexor has settled down (approximately 10 ms), the begin conversion command is issued. When the analog-to-digital converter completes its cycle, the value is transmitted back to the recording system using the other coaxial cable.

The analog-to-digital converters have a resolution of 12 bits (1 part in 4096), and they use a sign/magnitude binary code. The 12 binary bits are grouped into four groups of three bits each and are transmitted as four octal digits in standard ASCII code. (The sign bit is transmitted as the most significant bit of the most significant digit, so that the actual digitized value is obtained by subtracting 4000 octal from the received magnitude.) Thus both the command characters and the reply are in standard printing graphics running at a standard speed (300 baud). After the 4 octal digits are transmitted, the control section adds the multiplexor address as two additional octal digits and a carriage return character. The multiplexor address allows the recording system to verify that the correct number of channel advance characters were received by the digitizer. The carriage return is added for compatibility with standard input routines that buffer input to a program until a carriage return character is received.

Although it has never been used, the control section also recognizes five other ASCII characters. These five codes could be used to remotely perform any function. These codes were intended to provide the ability to change the gain of the electronics and to re-zero the system, but neither capability has proved to be necessary.

TRANSMISSION MULTIPLEXOR

The cables from every hole at a site terminate in one transmission multiplexor. The function of this circuit is to allow a large number of holes to be controlled by a single RS-232C port driven by a telephone line or by a micro-computer.

The transmission multiplexor accepts one RS-232C input from a controller and logically connects this input to any one of eight coaxial cable pairs. (The connection also includes level translation between standard bi-polar RS-232C levels and the TTL standard levels used for the transmission to and from the holes.) The transmission multiplexor recognizes the eight ASCII control codes with octal values 20 through 27 inclusive. When one of these codes is received, all subsequent characters are transmitted to the hole identifed by the low order digit of the code, and all signals received from that hole are transmitted back to the control input. Thus if octal 24 is received, all subsequent characters are sent to hole number four, and all responses by hole number four are returned to the controller. The connection remains unchanged until the next control code is received. Note that both command characters and valid response characters can never include any of these address codes and the two can therefore be sent intermixed in any order. This greatly simplifies the switching system since all switches can be active all of the time.

DATA RECORDING

The data from all of the tiltmeters are sent to the campus of the University of Colorado in Boulder. The existing stations and several proposed new stations are shown in fig. 3.

We use two types of transmission from the tiltmeter sites to the University.

The first type of system uses a dedicated voice-grade telephone line between the central computer and the tiltmeter site. This type of system is used to link the NBS site with the recording system, for example. The transmission multiplexor is driven from the central computer via modems at each end of the dedicated voice-grade telephone line.

The central computer keeps time by counting the zero crossings of the 60 Hz power line. Every six minutes the computer sends the appropriate codes to activate the transmission multiplexor at the distant end and to sample the data from every active tiltmeter.

The computer can make several rudimentary checks on the data. The computer verifies that the reply characters are valid octal digits and that the channel address returned by the digitizer corresponds to the number of channel advance characters sent. The computer also checks that only seven characters were received (four digits corresponding to the digitized tilt, two for channel address and the carriage return), and that the characters were received promptly (delay of less than five seconds). If any of these checks fails, the computer retries the entire process up to five times. If the error persists, the reading is flagged as having a fatal error, and the operation continues with the next active channel. This process is repeated for every site linked in this way. The average time to record the data from a single channel is about one second. During this time all other tasks on the computer are stopped.

In order to drive this system, the computer must maintain a table of active lines, and a list of how many instruments are active on every line. This is done in a manner described below under software.

This system is extremely simple and requires very little hardware at each site. A failure of the central computer, however, stops the data acquisition, and recovery of the missed data is impossible, since the site controllers have no memory to store data and no ability to initiate a measurement. Although this is a weakness in principle, it has not made much of a difference in practice, since the digital recording system has been one of the most reliable parts of the system.

Note that in this system a given tiltmeter is identified by a series of hardware-related parameters: the hardware port to which a given dedicated line is connected, the transmission multiplexor port to which a given hole is connected, and the analog multiplexor input to which a given tiltmeter is connected. Each one of these address parameters is distinct, and different sites do not share any hardware.

We now have two sites connected in this way (one at NBS and one in the basement of JILA itself), and the software as written can support seven such sites.

The second type of recording system uses a time-multiplexed radio link to transmit the data to the University. This type of system is used to link the Erie site, for example.

In this system, the transmission multiplexor at the site is driven by a micro-computer at the site. The micro-computer is programmed in FORTRAN and the program code is stored in programmable read-only-memory (EPROM). The EPROM chips can be erased by exposing them to intense ultraviolet light for approximately 20 minutes, and can then be reprogrammed.

The program development is done on a small micro-computer in the laboratory having the same hardware configuration as the field device. The program development machine also has a floppy disk system which holds the FORTRAN compiler, the editor and the program text. The micro-computers are built around the Z80 central-processor chip and the S-100 bus architecture; this configuration is supported by many vendors and was chosen primarily because the components are available in our stockroom.

The micro-computer acquires the data in much the same way as the central computer does. The program keeps track of the time of day by reading a quartz crystal clock connected to the main computer bus. The table of active tiltmeters is stored in the EPROM along with the program code and cannot be changed in the field. Every six minutes, the micro-computer activates the multiplexor and reads the data from all of the active tiltmeters. It then performs essentially the same checks as the central computer does: if the reading is valid it is stored in MOS memory. If the reading is invalid after five re-tries, the value 8000 is stored instead of the data in error. Note that the digit 8 is not octal and is used to indicate a hardware error. It cannot appear in a valid transmission. (The eight must appear as the first digit of a value to be recognized as the hardware-error flag. In any other position it is treated as a transmission error. The three digits following the eight are reserved for specifying the type of error, but only the value 000 is currently used.)

This process continues for 60 minutes (ten data scans). At the end of an hour, the micro-computer activates the VHF radio link to transmit the data to the central computer in Boulder.

The radio link is composed of commercial VHF equipment. The radiated power is 25 w into a quarter-wave vertical whip mounted over a ground-plane on a mast approximately 8 m above the ground. A commercial frequency-shift-key modem of the Bell system 103 type is connected to the audio input of the VHF transmitter using a small transformer to match the output impedance of the modem (600 ohms) to the input impedance of the audio stage (50k ohms). An adjustable attenuator is also provided to avoid over-driving the audio stages. The data input to the modem is connected to an auxilliary RS-232C port on the micro-computer. The transmitter is switched from standby to transmit by means of a small reed relay activated by setting the request-to-send (RTS) bit in the RS-232C interface.

After switching on the carrier, the program waits ten seconds for the carrier to stabilize and then begins by transmitting the station preamble. This consists of two lines of text in ASCII code as follows:

xxxxx-start <carriage return><line feed>

yyddd hh:mm:ss <carriage return><line feed>

where xxxxx is the name of the station (e.g. ERIE), yyddd is the year and the day (e.g. 81182) and hh:mm:ss is the time of the micro-computer clock in 24-hour format and in universal time (UTC).

After the preamble, the micro-computer waits one second and then begins transmitting the data in an encoded form (to be described below in the section on software). When all of the data have been transmitted, the micro-computer sends a <carriage return>, a <line feed> and the post-amble consisting of the phrase

xxxxx-end <carriage return><line feed>

where xxxxx is the station identifier as above. The micro-computer then waits five seconds and turns off the carrier, returning the transmitter to the standby condition.

After the transmission, the micro-computer examines the state of the data buffer and computes if enough room exists for a complete hour's worth of additional data (10 scans X 6 characters/scan X the number of active channels). If the buffer can hold that many more characters, nothing is done and the new data is simply added onto the end of the existing data. The next transmission will thus repeat some old data along with the new data to give the receiving station a second chance in case the receiving station was down or the first transmission was received in error.

If the buffer cannot hold another hour's worth of data, the data in the buffer are pushed down by the number of characters recorded per hour, thus erasing the oldest hour of data. Thus each transmission starts by repeating old data and ends with the most recent data.

The data buffer can hold 2500 characters. This is approximately five hours of data if eight channels are active, so that every datum will be sent about five times. Alternatively, the central site can be down for about five hours with no loss of data.

Since the maximum number of characters transmitted is about 2600, the maximum transmission time is about 85 seconds at 300 baud.

If the power fails, the data in the MOS memory is lost, but the clock keeps running, powered by a small 9 v battery. When the power is restored, the program is automatically restarted. Since the entire program text (including the active tiltmeter table) is stored in read-only memory, the system resumes synchronously when the next scan time arrives. Only the data from the last transmission to the power failure is lost, and this will always be an integral number of samples. (While the power is off the data would be lost in any case since the tiltmeters themselves go off.)

The data is received at the University by a commercial VHF receiver connected to a quarter-wave vertical whip antenna on a ground plane mounted on the roof of the laboratory building (about 10 m above the ground). The audio output is matched into a frequency-shift-key modem by means of an impedance-matching transformer and a resistive attenuator. The output of the modem is connected to one of the RS-232C inputs of the central computer.

Although only one site is connected in this way, the system has been designed to support up to ten sites. The receiving hardware is shared by all of the sites. Each site is assigned a unique site identifier (stored in its EPROM) and a unique transmission time. These times will be two minutes after any even tenth of an hour (i.e. at 2,8,14, ..., 56 minutes after the hour). The central computer begins listening for a transmission at one minute after every even tenth of an hour (i.e. at 1,7,13, ..., 55 minutes after the hour). If a transmission is not expected, the task exits immediately. If a transmission is expected (see the software section below), but none is received by twenty seconds before the next even tenth, the computer stops listening and reports a time-out. (For example, if a transmission is expected at two minutes after the hour, the computer will listen from 1:00 to 5:40).

In this way up to ten sites can share the same frequency and the same receiving hardware. The stations are identified by their transmission times and by their unique station identifiers transmitted in the preamble and postamble.

These times were chosen so that the transmission would occur while the entire system would be idling between data acquisition cycles, and so that the entire transmission cycle could be completed before the next acquisition cycle by either the micro-computer in the field or by the central machine. This scheme also allows the various clocks to drift by up to one minute without loss of data.

Although this time-shared scheme has been implemented using VHF links, the identical scheme could be implemented using automatic dialers driven by the field .cro-computers. The protocol would be identical except that the micro-computer would dial the central site using a commercial direct dial telephone. After the connection was established, the protocol could proceed as in the VHF system.

We anticipate using the telphone version of this protocol to transmit the data from our new sites in Wyoming and Utah.

SOFTWARE FUNDAMENTALS

The acquisition and analysis software is written almost totally in standard FORTRAN. Although much of the software is machine independent, the command interpreter and the file system interface, although written mostly in FORTRAN, are dependent on the specific characteristics of the hardware and operating system we have (PDP 11/34 running RSX11M).

The software consists of a series of indepedent programs which can be run in any order. All of the programs read and write data in a standard format, termed SDF format. The entire analysis package is therefore characterized only by the standard format chosen for the data files. Any task of any complexity can be added: the only requirement on the author is that he conform to the standard file structure for files that are intended to be processed by other components of the package (internal scratch files are, of course, exempt from this requirement). This scheme allows every task to have access to commonly used utilities (e.g. plot the results, list the results, compute the spectrum of the results, etc.) but yet does not require the author of every task to write (or even to understand) the other components of the system.

The file structure used for SDF format differs in two important respects from previous time-series analysis packages (a.g. BOMM):

1. The files are always read and written in physical disk blocks

(512 characters). FORTRAN and the operating system provide direct random access to any file at this level, and this is the fastest way of performing disk input/output. BOMM (and many other packages) read and write data using FORTRAN binary input/output. This access mode requires that the data be read and written sequentially, and changing any value in the middle of a file requires the entire file from that point to the end to be re-written.

2. The first physical block of every file is a header block containing information about the rest of the file. If the file is a time series, the start time, time interval and number of values are stored here; if the file is a spectrum, the start frequency, frequency increment and number of values are stored here, etc. BOMM (and many other packages) store this information in separate files or do not store it at all, relying on the memory and/or on the notebooks of the investigator to keep track of what is what. It was our intention to eliminate the need for human-kept notebooks as much as possible. (Note that the "station-tape" format used at IGPP to store the data from the IDA network is based on the same general idea, but that much of this sort of information is stripped from the data files when the data is read from tape to disk.)

A very important parameter stored in the header block is the format value. This parameter is a small integer that defines the structure of the rest of the file. For example a format 0 file is a raw data file. It is the sort of file used by the data acquisition task to record the raw data. Every block of the file after the first is composed of 256 16-bit integers. There is one such file for every port active on the system, the file containing the data from port j being named TT<j>DTA.SDF. These files contain all of the data from that port in a merged sequence (i.e. if five channels are active on port 1, file TT1DTA.SDF contains five values for each sample-time stored sequentially by channel number and then by time).

As a second example, if the format parameter is equal to 2, the file is a single component time series. Every block of the file after the first is composed of 128 32-bit floating point numbers.

Every task checks to make sure that the requested operation makes sense on the requested file. It makes no sense to compute the spectrum of a format 0 file, for example, or to store raw data in a format 2 file. If two format 2 files are joined end to end, a diagnostic is provided if the start time of the second file does not follow logically on the end time of the first, or if the time intervals of the two time series are not the same.

In addition to modifiers that are specific to a given operation, all tasks will accept the file modifiers /st:j and /sp:k implying that the requested operation is only to be performed on data values between index j and index k inclusive. Thus any operation can be easily confined to any desired subset of the file. The constants j and k may be any values consistent with the length of the file. If the requested subset begins or ends in the middle of a disk block, the required packing and unpacking is done internally and is totally transparent to the user. This is a great simplification over the BOMM system in which a subset of a file had to be explicitly extracted and stored in a second file before it could be used.

There are two limitations inherent in the SDF format files. The most serious limitation arises from the fact that integers are stored as 16 bit quantities, and the maximum value for an integer is therefore 32767. Since subscripts are integer quantities, it is difficult to deal with time series longer that 32767 terms long. This limitation is not as academic as it sounds. At ten samples per hour, a time series can be no longer than 3276 hours long, and we already have several time series longer than this.

There is no easy solution to this problem, since it is fundamental to the hardware configuration we have. A partial solution to the problem has been implemented by defining the concept of a 'chain' file, i.e. a time series which is logically continuous but which spans several physical files. Many of the SDF tasks do not process the 'chain' automatically, however, so that the solution is rather awkward.

A second, less serious, limitation arises from the fact that the SDF system uses the RSX11M message handler to parse command lines typed by the operator. This relieves the tasks of the need to write a command parser, but does not allow floating point values to be passed in commands. The command parser converts numerical parameters to 16 bit quantities, so that here too parameters must be integers in the range -32768 to + 32767. A solution for this problem has been found, but it has not yet been implemented.

DATA ACQUISITION SOFTWARE

The data acquisition tasks run as standard tasks within the operating system. They are fully compatible with the SDF format. The tasks are scheduled by the RSX11M executive to run every six minutes.

The hard-wired sites are controlled by a task named DATLOG and the VHF-linked sites are controlled by a task named VHFLOG.

The task DATLOG is initiated by the executive every six minutes starting from the first even tenth of an hour following the system deadstart; VHFLOG is initiated one minute later and also runs every six minutes. Both tasks reference a system-wide common block named SDFCOM to find which ports are active at any time. (A system-wide common block is a dedicated area of physical memory defined when the system is generated. A task with the requisite permission may access this block by name by including a statement of the form

COMMON/SDFCOM/ variable list

at the beginning of a program.)

This common block is 64 words long. The first two words are the number of seconds since midnight when the task VHFLOG was last active in standard floating point format, and the next two words are the corresponding parameter for the task DATLOG. These words permit the respective tasks to detect a crash within the executive, since at every initiation the difference between the current time (in seconds after midnight) and the stored time should be 360 (neglecting the wrap-around at midnight which must be handled specially).

After these four words of time, there are seven words for each of the possible seven hard-wired sites and ten words for each of the possible VHF linked sites. (The extra words in the common block are reserved for expansion.)

If a given site is active, the corresponding word is made non-zero, either by direct entry from the front panel or by a small utility program. Furthermore, if a given site has less than sixteen active channels (all current sites fall into this category), there is one bit set in the word for each active instrument at the site. The least significant bit corresponds to channel one at the site, and the most significant bit corresponds to channel sixteen. For example if the first four channels are active at a given site, the corresponding common block word will be set to 17 octal. Each task scans its respective words until a non-zero entry is found.

In the case of the task DATLOG (for the hard-wired sites) every non-zero word represents a port which must be sampled every six minutes. If word j is non-zero, for example, the task must scan port j. It opens file TT<j>DTA.SDF and reads the header block. The header block of a format 0 file contains the number of active channels and the complete address of every active channel (i.e. the multiplexor code and the datalogger channel number). This information allows the task to activate the remote dataloggers and to record the data. The header block also contains the time corresponding to the first scan in the file and the number of scans currently in the file. This information is used to decide if the current scan follows logically after the last recorded scan. If it does, the data is added to the file, the header block is updated appropriately, and the task continues by searching for the next non-zero word in the common area. When all active channels have been serviced, the task exits.

If a hardware error occurs (e.g. a particular channel consistently reports a value containing a non-octal digit or does not respond at all), a value of zero is stored for that channel, and a message is entered in a special file (see below). If all of the channels on a given hardware port do not respond, a hardware failure in the modems or in the telephone line is probably indicated. The task places an appropriate message in the special file and goes on to the next active port without updating the header file (as if the task simply had not been active).

If the current scan does not follow logically after the last recorded scan (because of a system crash or a previous hardware error), the task pads the file with the requisite number of missing values by repeating the last value as many times as is necessary to make the current value logically follow the new last value.

The task reports all errors to a special file called DTALOG.DAY.

These entries give the date and time of the error, the number of repetitions attempted, etc. It is an ASCII text file and can be examined with the editor or listed on a terminal.

The task VHFLOG is somewhat different. It too examines its common area to find a non-zero word. However, if word j is non-zero, it is taken to mean that a transmission is expected in time slot j (i.e. if word one is non-zero, a transmission is expected at two minutes after the hour). If the word corresponding to the current time is zero, the task simply exits.

If the word corresponding to the current time is non-zero, the task goes into a wait state waiting for input from port 30 (the port connected to the VHF receiver). If no transmission is received in the time-out period, the task issues a message to a special file (named VHFLOG.DAY) and exits.

If a transmission is received, the task opens file VHF30<j>.SDF and compares the station identifer in the preamble with the code stored in the header block of the file. If the two disagree, an error is flagged in the message file and the task exits. If the two agree, the task compares the current time with the time reported in the preamble, and issues a message if the micro-computer clock is in error by more than one minute. It then receives, unpacks and checks the data (see below) until the postamble is received. At this point the task computes the first 'new' scan received (much of the data has usually been received already) and begins adding the data to the end of the file. If the first scan received does not follow logically after the last scan recorded, some data has been lost due to a failure. The task bridges this gap by repeating the first scan received as many times as is necessary to make the data file logically contiguous. (Note that DATLOG bridges a gap by repeating the last old scan while VHFLOG bridges a gap by repeating the first new scan. This is done only for programming simplicity in each case.) After storing the new data, the task updates the header file appropriately and exits.

The data transmitted over the VHF link is encoded to permit the detection and correction of transmission errors. The encoding scheme was developed within the following constraints:

1. Although the computers all transmit and receive 8-bit ASCII codes, in fact it is awkward to transmit the 8-th bit to a user program in RSX11M, and to do so introduces other awkwardnesses. The coding scheme must therefore use only 7-bit ASCII with the 8-th bit ignored (always set to zero is the RSX11M convention).

2. The operating system uses some of the ASCII control codes for a simple line editor (e.g. entering control-U erases the current line, etc.). This use of control characters can also be by-passed, but to do so again introduces awkwardnesses. Thus a coding scheme should not use control characters, but should limit itself as much as possible to printing graphics.

3. The data as transmitted from the hardware at the tiltmeters

are fundamentally octal digits, and it is simplest to preserve this fact in the transmission. The encoding scheme should therefore map the eight octal digits into ASCII characters in such a way that the resultant characters are as 'orthogonal' as possible, that is, that the eight characters differ by as many bits as possible so that an error is unlikely to turn one legal character into another one.

4. It is desirable to be able to incorporate an indication of a hardware failure into the transmission format. The micro-computer uses a data value of 8000 to signify a hardware failure, and it would be desirable to be able to continue this convention in the transmission scheme so that the central site could be made aware of hardware problems. It is also desirable to have more than eight codes available for channel identifiers, since in general more than eight channels may be active at a given site.

The codes meeting these requirements are listed in table 1. In addition to all of the conditions listed above, these codes also have the property that (at least for the octal data codes) the low three bits of the code are exactly the value being transmitted. (This is not true for the extended address codes, all of which are ASCII control codes. These codes are chosen because they have no special significance to the operating system and are passed to a user task even though they violate criterion two above. These extended codes are only valid in the channel address field.)

The micro-computer encodes each data value to be transmitted into a seven character string in the following format:

xxxxyzz

where x is the translation of the appropriate octal digit in the data value, y is the channel number (using the extended values for channels after number seven), and z is the checksum obtained by summing the lowest three bits of the first five characters and adding this result to octal 40. The resulting checksum is always a printing character having an octal value ranging from 40 to 103. Note that the checksum is sent twice for re-synchronizaiton if a character is lost.

The task VHFLOG reverses the process. It decodes the first four characters using the following scheme:

1.If the value of x is exactly given by a value in table 1, report that value for that digit, set the error flag to zero and proceed with the next digit.

2. If the value of x is not given in table 1, but a value in table 1 differs from x by one bit, use that value for the digit, set a one bit error and return. Table 1 is constructed so that one bit errors are uniquely correctable.

3. If the value of x differs by more than one bit from any value in table 1, find the value in table 1 with the smallest number of bits different. This correction is not unique in general. Report this value for the digit, set the multiple bit error flag and proceed.

When the four x digits have been processed, decode the channel identifier using the same algorithm and allowing the extended codes as well.

Then compute the checksum and compare it with the transmitted value. If all is well, proceed. If a checksum error is indicated, go back and find the character with a multiple bit error and correct it using the indicated checksum if to do so would not raise the level of the multiple bit error. If this process does not yield a consistent value, flag a fatal error and exit, hoping for better luck next time.

If a checksum error is detected and the two checksum characters are different, it is possible that a character has been lost. In this case the program searches the characters in the buffer in the forward direction, looking for a valid scan pattern (i.e. a string of characters in the format outlined above). If such a pattern is found, the program re-synchronizes the decoding process with the valid pattern. If this process projected backwards to the current position implies an incorrect number of characters, the error is corrected by inserting the fewest number of characters at the current position to restore the synchronization.

This decoding process continues until the postamble is reached. If at that point the decoding process is in the middle of a scan, then an unknown number of data values have either been lost in transmission or added in the patching process. The program attempts to decide which of these two possibilities has occurred by examining that portion of the transmission that overlaps with data already received, but this process may not be successful; if no obvious correction can be found the program exits and waits for the next transmission.

This decoding algorithm is based on three assumptions, all of which are borne out by our experience:

1. The overwhelming majority of transmission errors are single bit errors, and these are always uniquely correctable by the code.

2. When an incorrect number of characters is found, it has almost always been the result of a lost character rather than an extraneously added one. If the start bit of a character is missed, that character will be appended to the one following it, with the result that the number of characters received will be too small. Since the transmission proceeds at the full theoretical bandwidth of the 300 baud channel, it is essentially impossible to receive additional characters (except possibly in the unimportant times when the carrier is switched on and off before the preamble or after the postamble).

3. Almost all transmissions are essentially error free, with at worst a small number of single bit errors. Serious multiple bit errors are usually encountered under adverse conditions (e.g. in the middle of a lightning storm), and the optimum strategy seems to be to exit if no obvious correction exists and wait for the next transmission.

PRELIMINARY DATA REDUCTION

Most of the experimental problems in installing and operating the tiltmeters were solved by September of 1980, and we have confined our analyses to the data acquired after October, 1980.

The first step in the data reduction is to remove the obvious errors produced by instrumental failures. The amount of data lost due to all hardware causes was quite small, and was confined to approximately eight relatively short intervals which occurred on the average of about once per month. (This total does not include a larger number of failures lasting only a few hours. These short outages were patched by interpolation.) Since the longer outages were relatively rare events, we decided to simply omit those periods from the analysis and not to try to fill in the gaps.

The data were lowpass filtered and then decimated to one sample per hour. The filters we used were symmetrical lead-lag filters and were constructed by applying a suitable windowing function to the idealized low-pass filter response function (i.e. sin x/x). The windowing function we used is the "Blackman" window (Oppenheim and Schafer, 1975). This window is basically a modified cosine taper. It has the desirable property of having the greatest stop-band attenuation of any of the commonly discussed windows. This stop-band attenuation is, of course, bought at the price of a wider transition zone between the pass-band and the stop-band. This wide transition zone is not too serious in our case since the tidal data is far below the Nyquist frequency, so that we can afford a wide transition zone and still not have an appreciable problem with aliasing.

The filters were constructed so as to have unity gain from d.c. to approximately 0.4 cycles/hour, and to fall to a transmission of approximately -80 db at 0.8 cycles/hour. These filters have a half-width of eleven terms.

Filters of this type are well suited to tidal analysis since they introduce no phase dispersion into the filtered data. They produce a time delay of an integral number of sample points (the delay is equal to the half-width of the filter function at the start and end of the time series); when this delay is accounted for by altering the start time and the end time appropriately, the resulting time series has a phase shift that is identically zero at all frequencies. The amplitude response in the pass-band is smooth so that calculated power spectra usually need not be corrected for the filter transmission function. These filters have the disadvantage that their stop-band attenuation is not very great and is not monotonic. This is not a serious problem provided that the power in the stop-band has no sharp features.

The decimation to one sample/hour reduced the number of data points by an order of magnitude, and made the subsequent task much easier. Since we are only interested in low frequencies, no information was lost by this process.

The next step in the process is to examine the power spectra of the data. The spectrum of the data from one of our horizontal pendulums running at the NBS site is shown in fig. 4. Several points are immediately apparent:

1. The power near 3.8 cycles/day is a direct measure of the non-linearity of the system. The largest single peak in the spectrum occurs at 1.93 cycles/day (the M2 tide); any non-linearity in the system must therefore generate a peak at 3.86 cycles/day. (A non-linearity whose dominant effect is to produce the third harmonic of the input power will also produce a signal here as a result of tripling the diurnal tides.) Since this is a relatively quiet part of the spectrum, this is a sensitive test for a non-linear response function. We conclude that any such non-linear response is too small to be measured. The spectrum may be used to set a lower bound on the non-linearity. If the response of an instrument is R when the applied tilt is T, then

R=k(T + qT**2 + cT**3)

where k is the gain of the tiltmeter and q and c are less than 0.005.

2. The power between the diurnal and semi-diurnal peaks (i.e. near 1.5 cycles/day) is a reasonable estimate of the ultimate precision attainable in any tidal analysis. The power at these frequencies is a measure of the broadband, low-frequency noise at a site; assuming that the instrument is not sensitive to the larger diurnal and semi-diurnal thermal disturbances, the power in this band is a good measure of the random noise that will limit the accuracy of a tidal admittance estimate. It is unlikely that power in this frequency band can be removed by fitting the tilt data with any simple function of the usual locally measured variables (e.g. local temperature, etc.), since the wavenumber spectrum is unknown in general. Using the observed spectrum, we find that the signal-to-noise ratio measured in this way is about 37 db, so that we should be able to determine the tidal amplitudes with an uncertainty due to random low-frequency noise of about 1%. The actual variance is larger than this (see below), so that we have not totally eliminated other diurnal and semi-diurnal effects (the most important of which is probably temperature).

TIDAL ANALYSIS

In order to compute the tidal admittance, we must compare the observed tilt tide with theory.

Our previous experience (Levine, 1976,1978) suggests that the ocean load may be quite different for the semi-diurnal and diurnal components, so that the fitting program must allow for at least four degrees of freedom (two amplitudes and two phases). We will not be able to get a good estimate of the long period admittance, however. Our data set is not long enough to begin with (only about eight months of data at most), and we must break the data set into several shorter pieces so as to get several estimates of the admittance in each frequency band.

In order to estimate the admittance, we have used an expansion of the potential in spherical harmonics (Munk and Cartwright, 1966). In this method, the potential at some time, t, is given by:

$$W(t) = \quad C \quad (t) Y \quad (\quad)$$

where W(t) is the tidal potential, C (t) is a time dependent function that must be calculated from the astronomy and Y ( ) is the normalized spherical harmonic of the station co-latitude, , and East longitude, . The sum is taken over all physically reasonabe values of n and m. The dominant contribution to the earth tide signal comes from the first non-vanishing term in the expansion, which has n=2. For n=2, the allowed values of m are 0, 1, and 2; the first is a long-period term, the second term is nominally diurnal and the third term is nominally semi-diurnal. The next order term, with n=3, is about one percent of the n=2 term at our latitude and may be ignored when estimating the diurnal and semi-diurnal admittances. (It is, of course, the dominant source of the three cycle per day tide.)

This method has many of the advantages of the Fourier decomposition methods, e.g. an expansion of the theoretical tides using the Cartwright-Tayler-Edden (1971,1973) potentials, without many of the disadvantages. This method allows the admittance to be different for the long-period, diurnal and semi-diurnal components. This is physically reasonable since both the ocean load and the correction for local topography are likely to be different among these widely disparate frequencies (Levine and Harrison, 1976; Levine, 1978). At the same time it does not allow different admittances for the components within one of these broad bands (the method can be used to allow structure within a band by incorporating leads and lags in the fitting process, but we have not used leads and lags with any great success in our previous work). This is a physically reasonable first hypothesis since appreciable structure within a band must be caused by a high-Q resonant structure, and such structures are unlikely. It is true that the same level of resolution could be achieved in the Fourier decomposition methods by summing the various diurnal components into one time series. However, the full Cartwright-Tayler-Edden (op. cit.) potential contains several hundred terms, and computing them all is quite a job. We feel that the full-blown Fourier decomposition is only justified if one has physical reason to believe that the tidal admittance is in fact modified by a high-Q resonance (e.g. in our previous efforts to estimate the modification of the diurnal admittance due to the well known core resonance as discussed in Levine, 1978).

At the same time, the method is far better than simply fitting the data with a theoretically generated tilt series. Such a fit has only one degree of freedom (i.e. an overall admittance), and this is both physically unreasonable and numerically dangerous. If the tidal data contains any drift, such a scheme will treat the drift as a long-period term unless the data set is very long indeed, and the resultant long-period admittance will appear to be much too large. Since this type of fit has only one degree of freedom, the resultant admittance is simply incorrect.

The actual calculation of the tidal admittance using this method proceeds in two parts:

1. The potential is calculated for each time at which a tilt tide value has been measured. The potential consists of five terms for each such time. The five terms correspond to the amplitude of the n=2, m=0 term, the real and imaginary parts of the n=2, m=1 term and the real and imaginary parts of the n=2, m=2 term. Each term is multiplied by its respective normalized spherical harmonic, and the resultant five component series is stored as a single SDF format file in which the header block specifies that the file is a set of five merged time series (except for the fact that the data are 32-bit real quantities rather than 16-bit integers, the format is the same as format 0 described above).

2. Once the potential file has been generated, a second task computes the least-squares coefficients between each time series component and the tilt data. Note that one potential file can be used to fit all of the data sets from a single site. The least squares parameters obtained in this way are in fact a function only of the azimuth of the tiltmeter and the effective Love numbers at the site, so that these numbers can be converted to absolute admittances. As a check on this procedure, the entire process is repeated with a theoretical time series.

This analysis has been done using several different lengths of data to investigate the stability of the admittance and its possible time dependence. It is unwise to use very long time series in the analysis, since only a few such calculations can be performed and only a poor estimate will be obtained for the fluctuations in the admittance. On the other hand if very short time series are used, the various components may not be well separated. We have chosen to use series 28 days long. By doing so we exploit the natural lunar-month periodicity of the tidal signal, and we can make eight estimates of the admittance using our data.

The most stable admittance should be that for the M2 component, since its frequency is far removed from the usual perturbations due to thermal effects. In table 2 we show the results of our estimates of this admittance. (In fact what is listed is not the M2 admittance, but the n=2, m=2 admittance, which contains all of the semi-diurnal components. However the M2 component is by far the largest contributor to the n=2, m=2 time series, and it is not unreasonable to refer to this term as the M2 admittance.) The data we have used were acquired using one of our horizontal pendulums in a hole 33 m deep at our NBS site. The azimuth of this pendulum is 232 degrees measured clockwise from North. The NBS site is located at 39.992 N., 105.269 W.

As can be seen from the table, the tidal admittance is surprisingly close to unity (a time series that corresponded exactly to theoretical expectations would have an admittance of unity and a phase shift of zero). Note that this calculation contains no adjustable constants. These admittance calculations make no correction for ocean loads, local topography or cavity effects, and inclusion of these effects may change our absolute admittances by a few percent. Nevertheless the agreement is encouraging.

In the present context the stability of the admittance is at least as important as its absolute value, since the stability is a direct check on the stability of the instrumental calibration. Although our preliminary analysis led us to expect random fluctuations in the admittance on the order of 1%, the observed fluctuations are approximately six times larger. This implies that the noise spectrum is not white, and that there is some residual sensitivity to extraneous effects (presumably temperature). The amplitude of the extraneous signal necessary to account for the observed fluctuation in the admittance is on the order of five nano-radians, and in view of the large surface-temperature cycle (20 C difference between day and night is not unusual) it is not unreasonable to assume that this is in fact a real tilt of the hill. (The instruments are located on the side of a rather steep hill. The slope of the hill is towards the South-East, so that the instrument-side of the hill is exposed to the sun during most of the day.) Our long-period analysis (see below), however suggests that these thermal effects are instrumental rather than real, and our main effort at the moment is finding out which part of the system is responsible.

Nevertheless, our tidal analysis leads to the following conclusions:

1. There is no secular change either in the calibration of the tiltmeter or in the elastic properties of the installation. Although the admittance is noisier than one would calculate on the basis of the white noise model, the admittance shows no secular trend.

2. The instrument and the site show no observable non-linear behaviour. The strongest piece of evidence for this is the absence of any excess power at 3.86 cycles/day.

3. The agreement between theory and experiment is surprisingly good. Although the site is poor in many ways, the observed admittances do not differ from what one would expect on the basis of classical tidal theory.

ANALYSIS OF LONG PERIOD EFFECTS

The second area of our analysis involves a study of the secular tilt records. We have no theoretical prediction of what these data should look like so that our analysis has been confined to the following three areas:

1. Are there any unreasonably large "tilt episodes?"

2. What is the effect of local temperature and rainfall?

3. What is the coherence among closely spaced tiltmeters?

In order to investigate these questions we have used the data from two instruments located at the NBS site in Boulder. The two instruments are a pair of vertical pendulums located in a hole 16 meters deep, and a pair of horizontal pendulums located in a hole 33 meters deep. The two holes are spaced a few meters apart.

The holes are located on the side of a hill at the Eastern edge of Green Mountain. We would expect data from such a site to be modified by the topographic anomaly and the crustal inhomogeneity, and it is not unreasonable to expect secular tilts. Our previous work using a laser strainmeter in the Poorman mine nearby showed secular strains on the order of one micro-strain per year (i.e. a secular increase in the length of the strainmeter baseline by about one part per million per year), and annual tilts at the micro-radian level are therefore not unreasonable. Tilt "episodes" showing a time derivative much higher than that implied by the overall annual rate or tilts much larger than a few micro-radians are suspect.

In fig. 5 we show the data from the four tilt sensors. There are no tilt episodes in any of the records (compare fig. 4, Wyatt and Berger, 1980). This negative result is encouraging. It re-confirms the hypothesis that carefully designed instruments, be they strainmeters or tiltmeters yield records that are totally consistent with simple elastic models of a site.

We attribute our smoother records to our better installation. Our instruments are deeper, and they are almost certainly more rigidly attached to the surrounding material than those at other sites (compare fig. 3, Wyatt and Berger, 1980).

It is almost certainly true, however, that the tilt records we have obtained are contaminated by other effects. To show this we have plotted the cumulative rainfall at the bottom of the figure. It is clear that both of the vertical pendulums are sensitive to rainfall, while the horizontal pendulums show no such sensitivity. We attribute this to the fact that the vertical pendulums are only 16 meters below the surface. This effect is shown more clearly in the subset of the data plotted in fig. 6.

It is not clear if the overall drift is real. The two vertical pendulums show highly correlated long term drifts; the same is true for the horizontal pendulums. If these signals represent real tilts, they imply that both tiltmeters are tilting along axes almost exactly midway between the two sensors; while this cannot be ruled out it seems somewhat implausible. We feel that it is more likely that the two sensors in each capsule are responding to some external event. (The event cannot be influencing the digitizing system. Its response has not changed to within one least count.)

It is not clear what this effect is. In fig. 7 we show a section of the same data sets plotted along with the local temperature. The situation here is quite confusing. Channel number one of each instrument is sensitive to temperature and channel two is not. It is unlikely that these signals represent real tilts, since channel one of the vertical pendulum is installed with its sensitive axis along an azimuth of 218 degrees (measured clockwise from North) while channel one of the horizontal pendulum has its sensitive axis along an azimuth of 142 degrees. These two azimuths are almost orthogonal. In each case, both of the channels are nominally identical in all respects. We have no model for this effect at this time; an examination of the temperature sensitivity of the system is our highest priority task at the moment.

We may summarize our long period results as follows:

1. An instrument placed in a hole 16 meters deep is definitely sensitive to rainfall. There is no question that instruments must be deeper than this. Our evidence to date suggests that 33 meters is deep enough.

2. Even in an instrument 33 meters deep, there is a residual temperature dependence on one channel that remains unexplained. Note that the second, nominally identical, channel shows no such effect. It is possible that this signal is generated by a true thermally-induced tilt, but we regard this possibility as unlikely. It would require that the instruments had one sensitive axis exactly aligned along the direction of the thermal tilt, and that the thermal tilt was in a different direction at the two depths. We regard it as more likely that some as yet undiscovered asymmetry in the instrument is responsible.

CONCLUSIONS

We have constructed a borehole tiltmeter that can operate for long periods of time at tidal sensitivity, and can yield tidal admittances in good agreement with theory. We have also developed ancillary recording and processing hardware to acquire the data from several such instruments installed at widely separated sites.

If the tiltmeters are installed at least 33 meters below the surface, the instruments do not show any effect due to local rainfall, but some residual temperature dependence is present in some cases. We attribute this sensitivity to some as yet undiscovered asymmetry in the system. We are currently acquiring data at the Erie site that we hope will answer these questions.

None of our instruments has shown a tilt 'episode' in approximately eight months of operation. All of our data to date is fully consistent with simple elastic models of the earth.

FUTURE EXPERIMENTS

Two of the major challenges in experimental tilt measurements are demonstrating long-term coherence between closely spaced tiltmeters and subsequently measuring long term crustal deformation, and, secondly, measuring tidal tilt with a sufficiently high signal-to-noise ratio and with relatively small or modelable cavity and topographic influences so that inversion for crustal material properties can be accomplished. The unique and fairly well understood geologic features of Yellowstone National Park make it an ideal location to study both the influence of geologic contrasts on tilt tides and to measure long term tilt associated with tectonic crustal deformation.

Yellowstone National Park has recently been the site of a large number and variety of geophysical and geological investigations which are summarized in Eaton, et al. (1975) and Smith and Christiansen (1980). The results of these investigations suggest the presence of a large, hot, and probably partially molten, body in the lower crust and upper mantle below the center of the park.

Smith and Christiansen (1980) have obtained a detailed three-dimensional density/slowness model to a depth of 100 km using results of the simultaneous inversion of complete Bouguer gravity data and teleseismic p-wave data. This model shows a widening with depth of the anomalous low-density and low-velocity body. The margins of the upper portion of the body are located below the boundary of the caldera. The shape of the anomalous body and the contrast in seismic velocities suggest similarities with the elastic dilatancy inclusion model of Beaumont and Berger (1974) which predicts up to 40% changes in tilt and strain amplitudes with a 10% reduction in Vp. Beaumont and Berger (1974) parameterized the numerical results in order to relate changes in elastic properties, and the associated Vp and Vs changes, to surface displacements, tilts, and strains calculated by the finite element models. From this they show that the tilt anomaly is a function of Poisson's ratio while the strain and displacement anomalies are functions of the areal bulk modulus.

Harrison and Flach (1976) have examined the connection between partial melting and variations in elastic parameters and p-wave velocity which might be expected for the Yellowstone anomalous body. Fig. 8 is a bulk modulus (K) - shear modulus (U) plot which shows the locus of points which correspond to a 10% reduction in Vp. Also plotted are trajectories which were computed assuming that the Vp reduction is due to partial melting using the Walsh (1969) theory for ellipsoidal melt inclusions. The important parameters effecting K and U are the aspect ratios (thickness/diameter ratios) of the melt inclusions, and the proportion of melted material. These parameters can be deduced by fixing the position of the Yellowstone body on the K-U plot using the Poisson's ratio obtained from the tilt anomaly inversion and some other elastic property (such as determined by Vp).

Finite element models were used by Harrison and Flach (1976) to compute the influence on the M2 tides for three cases corresponding to a 10% reduction in Vp. (I) A body extending to 100 km with material properties corresponding to very flat inclusions; (II) the same body with round inclusions and, (III) a body with the same material properties as (I) but extending to 200 km depth. The results for anomalous uplift (and fractional change in the gravity tide) and for tilt (and fractional change in the tilt tide) are shown in fig. 9. The effect on the gravity tide is small in all cases. The tilt anomaly is up to 40% for models (I) and (III), but is considerably smaller for model (II), for which there is little contrast in Poisson's ratio. The anomaly for model (III) is significantly more widespread than that for model (I), reflecting the greater extent in depth of the model (I) body.

Beaumont (1978) has further explored potential sources of spatial and temporal variations in tidal amplitude by considering the effects of stress hysteresis in dilatant zones. He suggests that a dilatant region can be modeled as a plastic-elastic crustal inclusion. He shows that near the inclusion the tidal admittance will depend on the relative tidal and tectonic stress rates and in addition will be anisotropic. He also predicts a nonlinear admittance when the tidal and tectonic stress rates are approximately equal. These effects will be of particular interest in Yellowstone since the anomalous body there is hot and possibly will behave as a plastic-elastic body. In addition, there is a reasonrbly high level of seismicity in and around the park. The largest earthquake had a magnitude of 7.1 and occurred in 1959 near Hebgen lake, just to the West of the park. The Yellowstone area is clearly subject to tectonic stresses, which, if they varied in time, would be reflected in time varying tidal tilt amplitudes.

An additional effect which may be observed at Yellowstone is a r lative phase shift between two tiltmeters, one of which is over the body and the other of which is not, due to the imperfectly elastic tidal response near the anomalous body. Zschau (1976) calculated that tilt tides are insensitive to a low-Q asthenosphere, but that load tide phase shifts of as high as several degrees may occur near ocean ridges and subduction zones. This effect has never been unambiguously separated from ocean load phase shifts due to the uncertainties in the modeli.g of the ocean tide. The effect on body tides of a vertical low-Q zone in the crust has not yet been calculated.

'ith several months of record, we can use our deep borehole ti.. ters to measure earth tides with uncertainties of less than 1%. The ffects of distant ocean load tilt can reach 20% in mid-(tinental areas (Harrison, 1976) but can be ignored when considering relative amplitude and phase changes over a comparatively small area such as Yellowstone. There are no significant problems with cav ty effects with vertical borehole tiltmeters, and topographic effects can be reduced by careful site selection and calculated with finite element models to within a few percent. The large expected amplitude anomalies of up to 40% should therefore be clearly observed. Time variations in tidal amplitude could be observed if the gain factors of the tiltmeters are stable.

In addition to tidal tilt studies, the Yellowstone National Park area is a good location to observe long term crustal tilting. Level lines established in 1923 and releveled in 1975-77 reveal a northeast trending elongated zone showing up to 700 mm vertical uplift (fig. 10). The area of uplift is coincident with the caldera rim and the two resurgent domes. The uplift rates of 11 mm/year are comparable to rates obtained in the volcanically active regions of Hawaii and Iceland (Smith and Christiansen, 1980). Using the steepest gradient, an average tilt rate of one microradian per year is obtained. The tilt rate is of course only the average over 50 years and is obtained from only two points in time. The deformation, if volcanic in origin, is likely to be episodic, and tilt rates over a short period of time could be much less or much greater. Leveling surveys are expensive and time consuming and therefore are not made very frequently. A microgravity benchmark network was established in the park in 1977, and will be re-surveyed every few years (Smith and Christiansen, 1980).

We propose installing tiltmeters at five sites (fig. 10) which straddle the zone of uplift and the margin of the essentially coincident anomalous body. These sites are preliminary and are subject to National Park Service approval. The proposed tiltmeter array offers a relatively inexpensive, convenient way of measuring crustal deformation at periods from hours to years. The coincident tilt, leveling and microgravity arrays will allow an evaluation of the consistency between point tilt measurements and large aperture geodetic measurements for long-period deformations. A sixth tiltmeter site is proposed near the Western entrance to the park. This site is in the vicinity of the 1959 Hebgen Lake earthquake. This magnitude 7.1 normal earthquake caused large vertical displacements (up to six meters) along exposed faults and down warping of Hebgen Lake by up to six meters (Smith and Christiansen, 1980). The area is still very active seismically and geodolite trilateration data shows continuing northeast-southwest extension, a pattern consistent with the sense of deformation of the 1959 earthquake and focal plane solutions from recent earthquakes (Prescott, et al., 1979).

REFERENCES

Beaumont, C. and J. Berger, 1974, Earthquake Prediction: Modification of the Earth Tilts and Strains by Dilatancy, Geophys. Jour. Roy. Astro. Soc., Vol. 39, pp. 111-21.

Beaumont, C., 1978, Linear and Non-Linear Interactions Between the Earth Tide and a Tectonically Stressed Earth, Proceedings of the Ninth GEOP Conference, Ohio State University, Department of Geodetic Science Report 280, pp.313-9.

Cartwright, D. E. and A. C. Edden, 1973, Corrected Tables of Tidal Harmonics, Geophys. J. R. Astr. Soc. 33, 253-264.

Cartwright, D. E. and R. J. Tayler, 1971, New Computations of the Tide-Generating Potential, Geophys. J. R. Astr. Soc. 23, 45-74.

Eaton, G. P., R. L. Christiansen, H. M. Iyer, A. M. Pitt, D. R. Mabey, H. R. Blank, Jr., I. Ziety and M. E. Gettings, 1975, Magma Benath Yellowstone National Park, Science, Vol. 188, pp. 787-96.

Harrison, J. C., 1976, Cavity and Topographic Effects in Tilt and Strain, J. G. R. 81, 319-328.

Harrison, J. C. and D. Flach, 1976, Eine Betrachtung der Erdgezeiten, 36th meeting of the German Geophysical Society, Bochum, April 1976.

Levine, J., 1978, Strain Tide Spectroscopy, Geophys. J. R. Astr. Soc. 54, 27-41.

Levine, J. and J. C. Harrison, 1976, Earth Tide Measurements in the Poorman Mine near Boulder, Colorado, J. G. R. 81, 2543-2555.

Munk, W. H. and D. E. Cartwright, 1966, Tidal Spectroscopy and Prediction, Proc. Roy. Soc. 259, 533-581.

Oppenheim, Alan V. and Ronald W. Schafer, 1975, Digital Signal Processing, Prentice-Hall, Inc. Englewood Cliffs, N. J., pp. 237-269.

Smith, R. B. and R. L. Christiansen, 1980, Yellowstone Park as a Window on the Earth's Interior, Scientific American, Vol. 242 pp. 104-17.

Prescott, W. H., J. C. Savage and W. T. Kinoshita, 1979, Strain Accumulation Rates in the Western United States Between 1970 and 1978, J. G. R. Vol. 84, pp. 5423-35.

Wyatt, F. and J. Berger, 1980, Investigations of Tilt Measurements Using Shallow Borehole Tiltmeters, J. G. R. Vol. 85, pp. 4351-4362.

Walsh, J. B., 1969, New Analysis of Attenuation in Partially Melted Rock, J. G. R. Vol. 74, pp. 4333-7.

Zschau, J., 1977, Phase Shifts of Tidal Sea Loading Deformations of the Earth's Surface Due to Low Viscosity Layers in the Interior, Proceedings of the 8th International Symposium on the Earthtides, Bonn.

Table 1--VHF Encoding System

VALUE	ASCII GRAPHIC	OCTAL VALUE	NOTES
0	X	130	1
1	9	071	
2	2	062	
3	S	123	
4	,	054	
5	M	115	
6	F	106	
7	'	047	
8	@	100	2
9	-	001	3
10	-	002	3
11	-	004	3
12	-	005	3
13	-	007	3
14	-	010	3
15	-	011	3

Note 1. The value of zero is never used as an address. The first channel is always number 1.

Note 2. This value is used for channel address eight and is also used within a data stream to signify a hardware error (see text). When used to signify a hardware error, it must be the first data digit.

Note 3. This value is only legal as a channel address. It will not be recognized if the program expects a data digit.

Table 2--M2 Admittance

HORIZONTAL PENDULUM

NBS Site (Boulder, Colorado)

28-day estimates

BLOCK NUMBER	TIME PERIOD	AMPLITUDE	PHASE(deg.)
1	Oct. 1980	1.06	-0.49
2	Nov. 1980	0.92	-7.60
3	Dec. 1980	1.01	-3.90
4	Jan. 1981	0.94	-4.50
5	Feb. 1981	1.03	-6.40
6	Mar. 1981	1.01	-3.10
7	Apr. 1981	0.92	-7.50
8	May 1981	0.99	-4.50
Avg.		0.984	-4.8
Std. Dev.		0.06	2.5

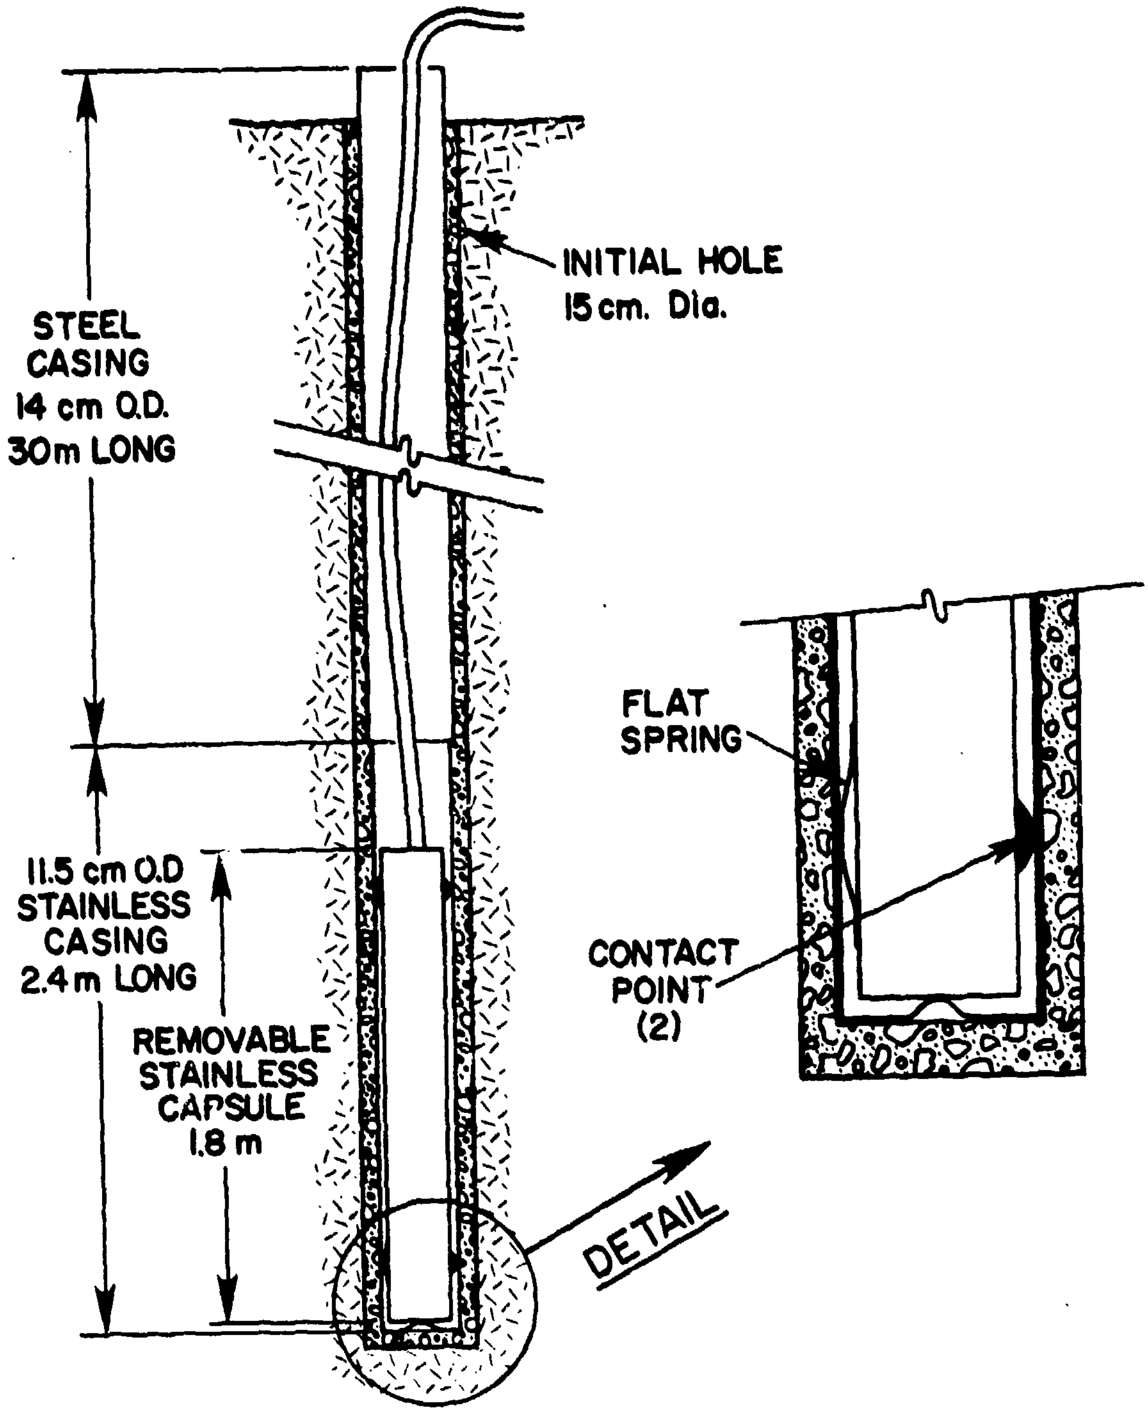

Fig. 1. A schematic diagram of the tiltmeter capsula installed at the bottom of a 33 meter deep borehole.

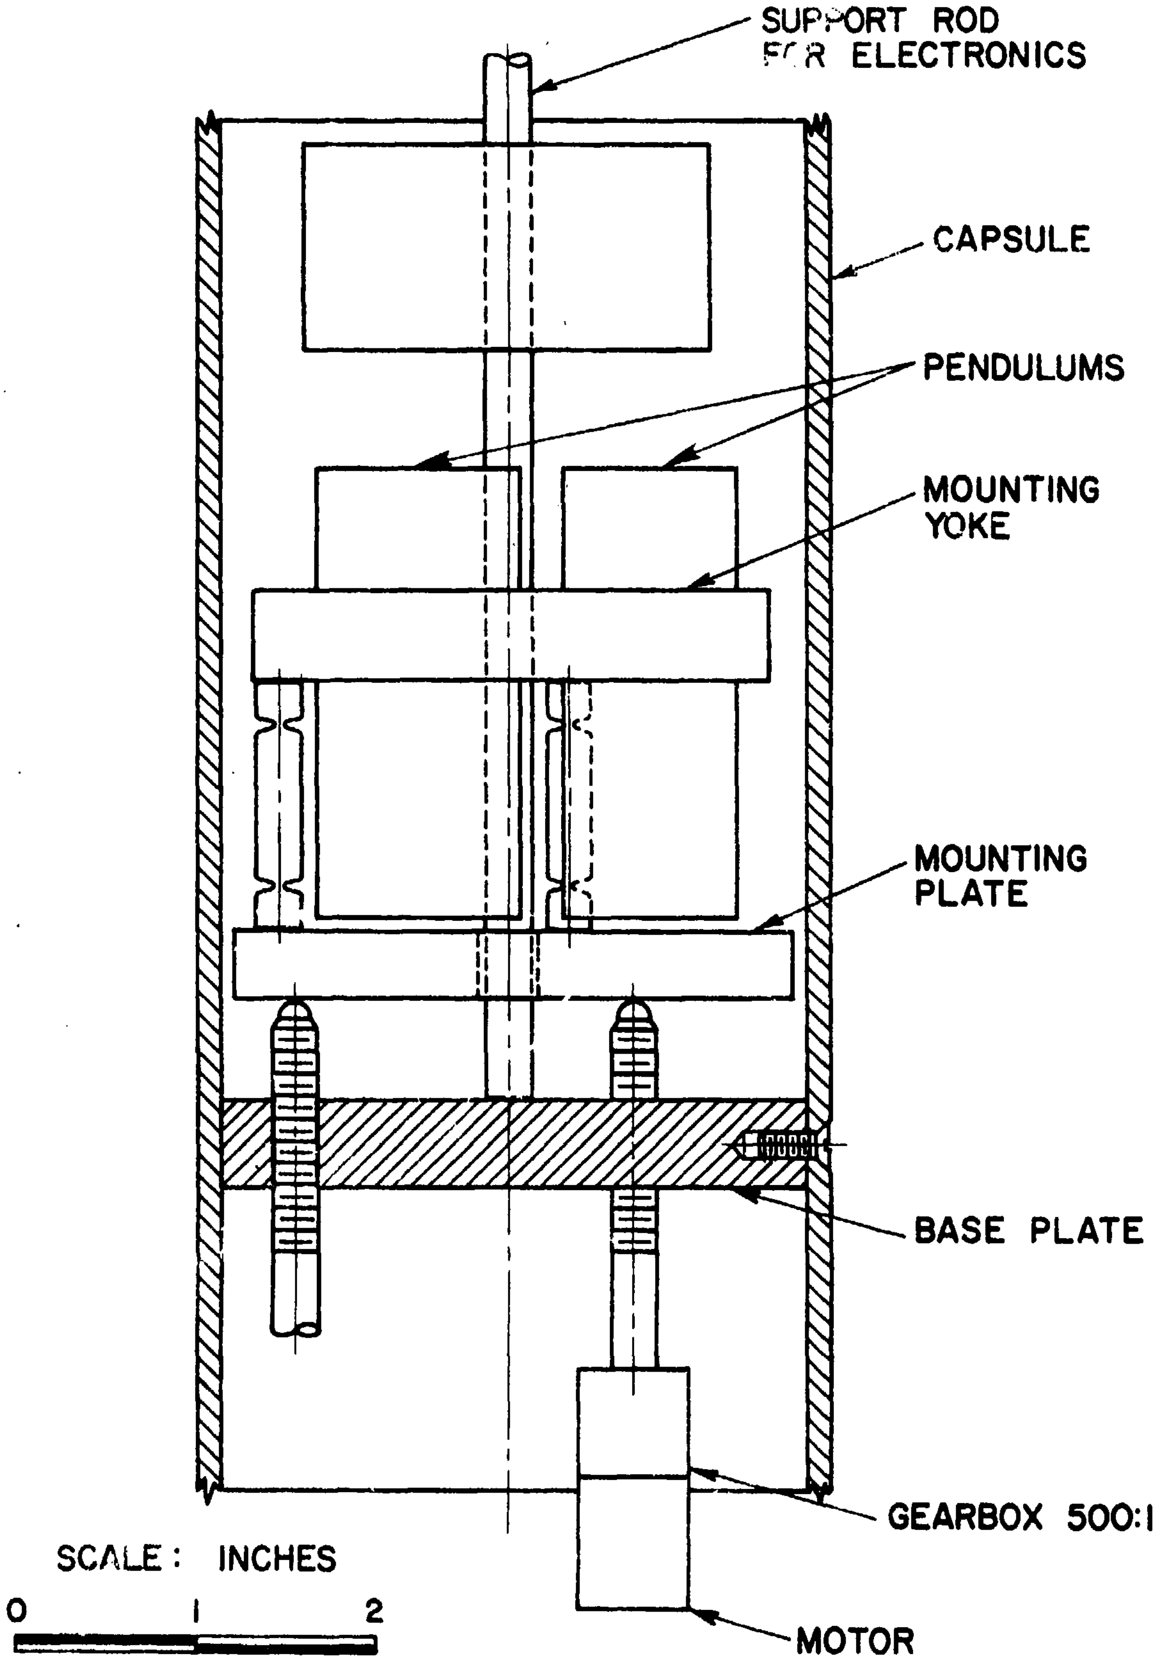

Fig. 2. Scale view of the tilt sensors mounted on the levelling platform inside of the capsule. The electronics package is above the tilt sensors inside the capsule and is not shown. Also not shown is the thermal insulation.

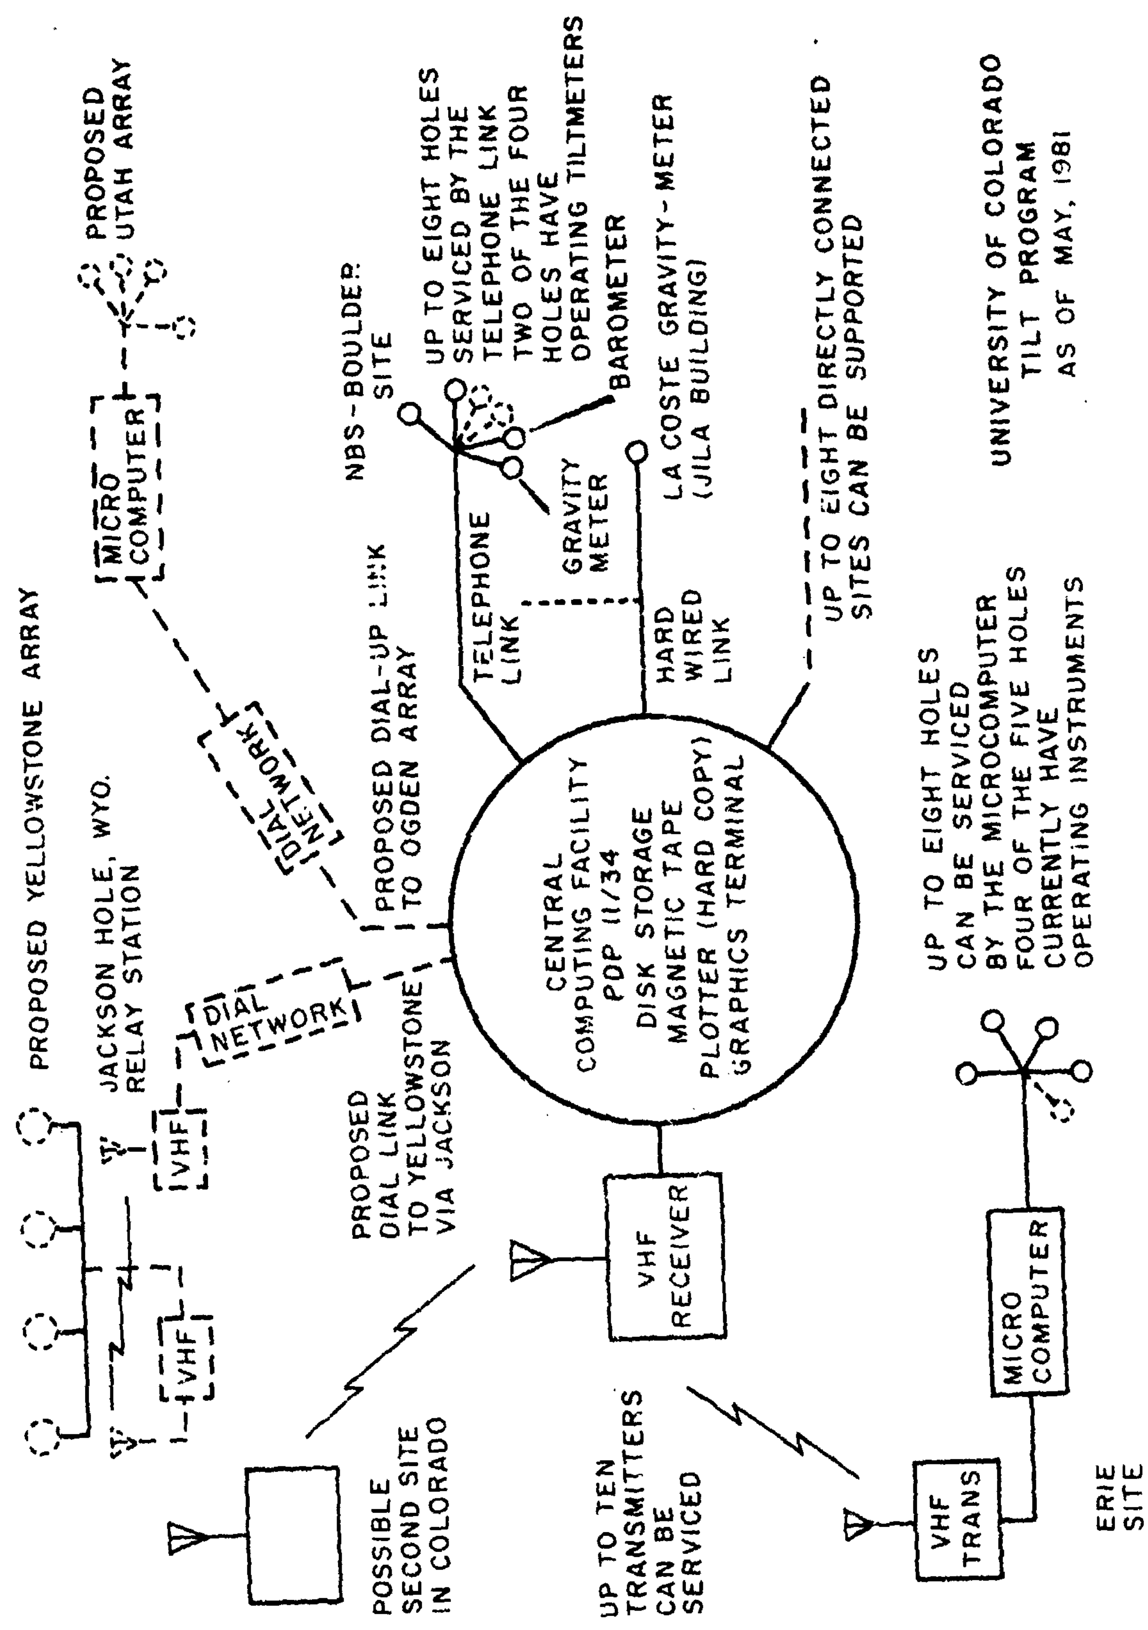

Fig. 3. Schematic diagram of the telemetry system. In addition to the two existing tilt sites at NBS in Boulder and at Erie, Colorado, we show the new sites proposed for Ogden, Utah and for Yellowstone National Park, Wyoming. These new sites will use a modified version of the existing VHF protocol as discussed in the text.

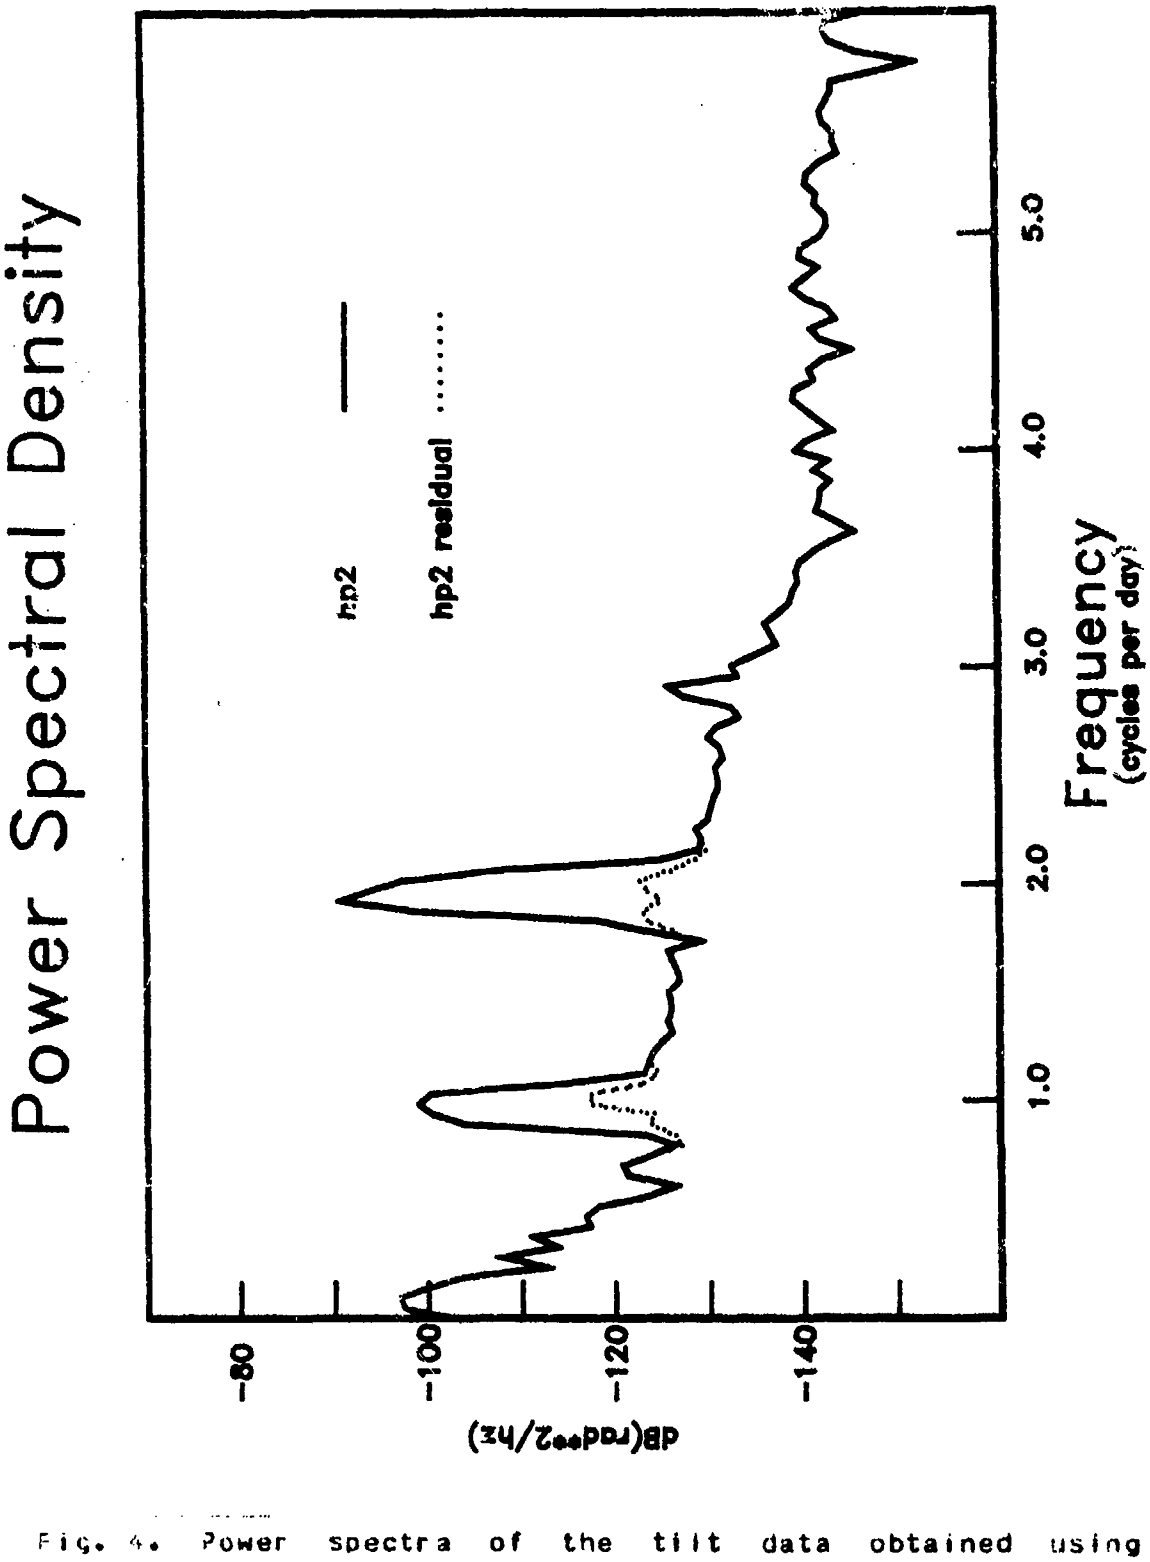

Fig. 4. Power spectra of the tilt data obtained using a horizontal pendulum at the NBS site in a hole 33 meters deep. The curve marked "residual" is the spectrum of the residual time series obtained by removing a least-squres fit of the tides to the data.

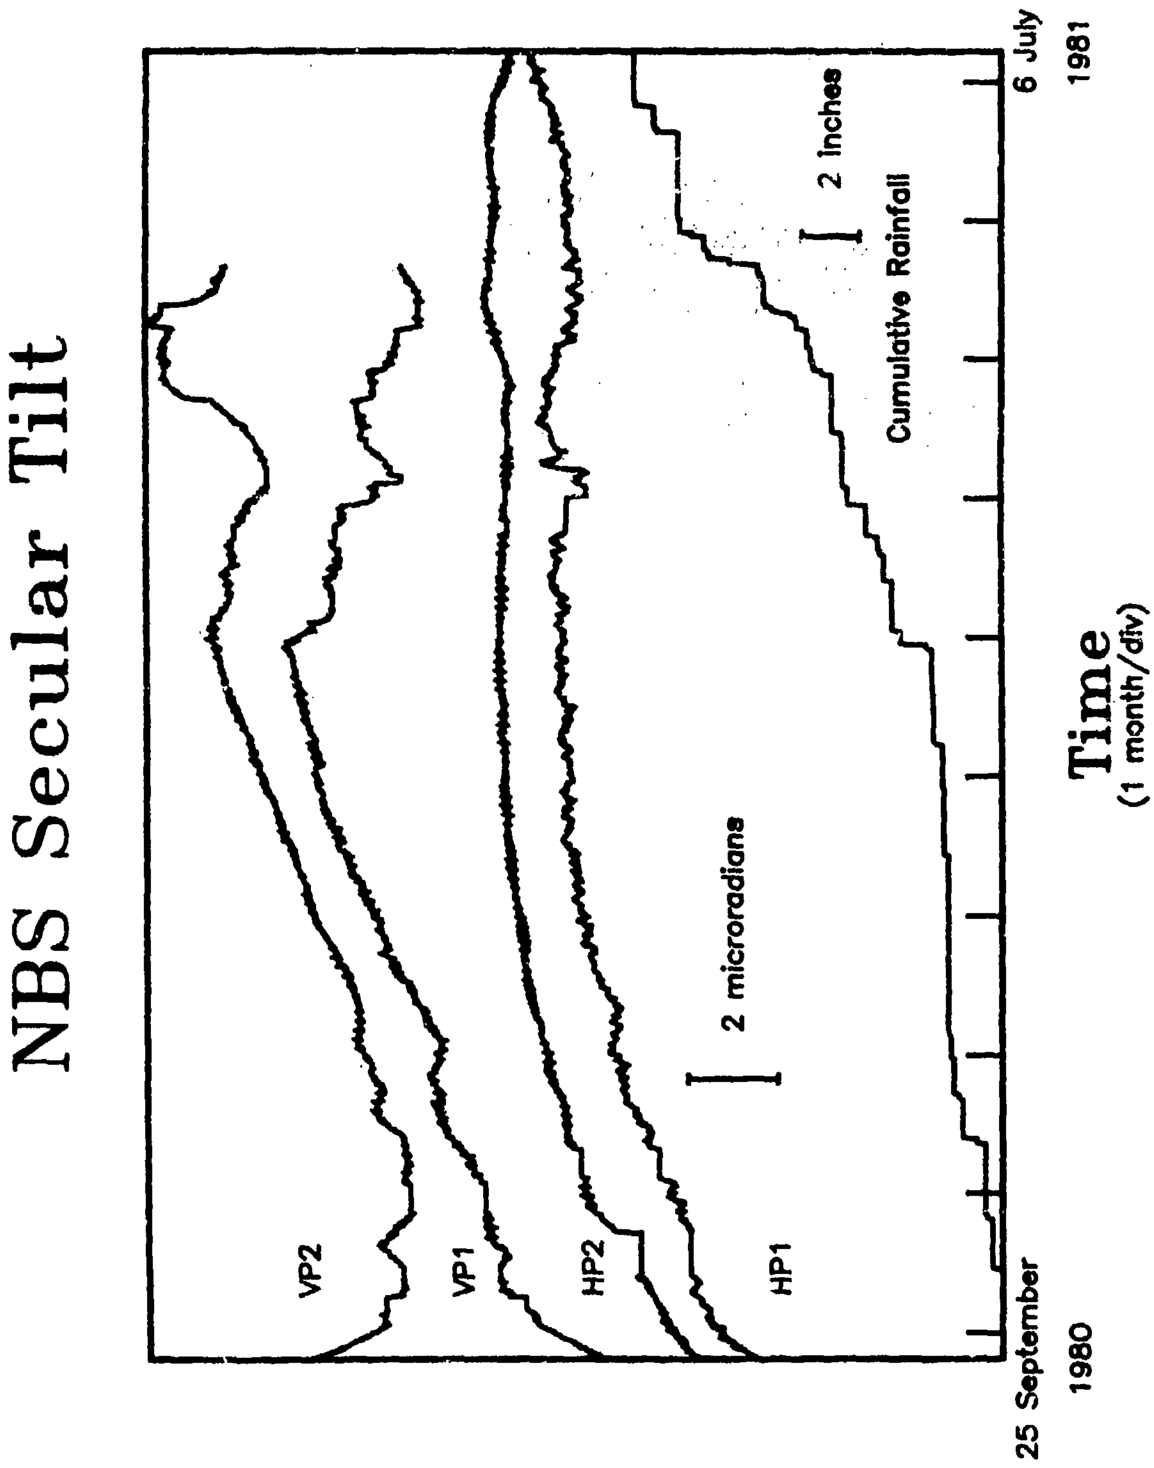

Fig. 5. Data from four tiltmeters at the NBS site. The traces marked VP1 and VP2 are data obtained using vertical pendulums in a hole 16 meters deep. The traces marked HP1 and HP2 are data obtained using horizontal pendulums in a hole 33 meters deep. The bottom trace shows cumulative rainfall taken from weather records. The vertical pendulums were moved to Erie late in May, 1981.

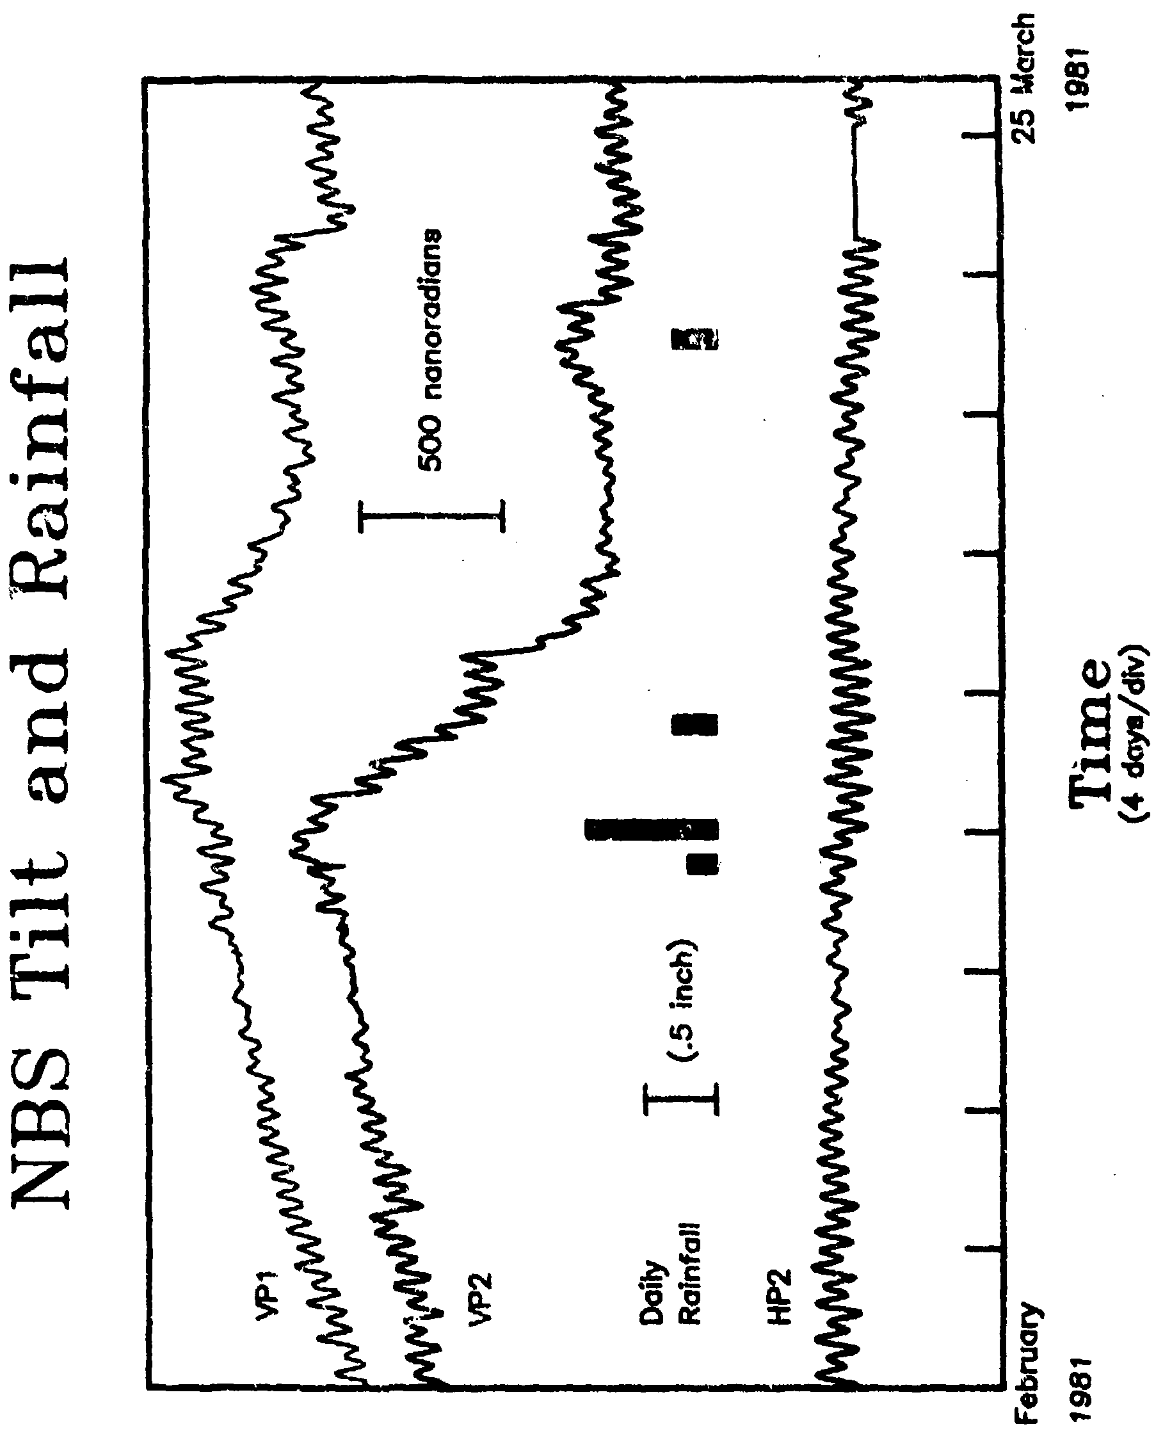

Fig. 6. A comparison of the sensitivities of different instruments to rainfall. The traces marked VP1 and VP2 are data from two vertical pendulums in a hole 16 meters deep. The trace marked HP2 is data from a horizontal pendulum in a hole 33 meters deep. Rainfall is shown by the black bars near the center of the figure.

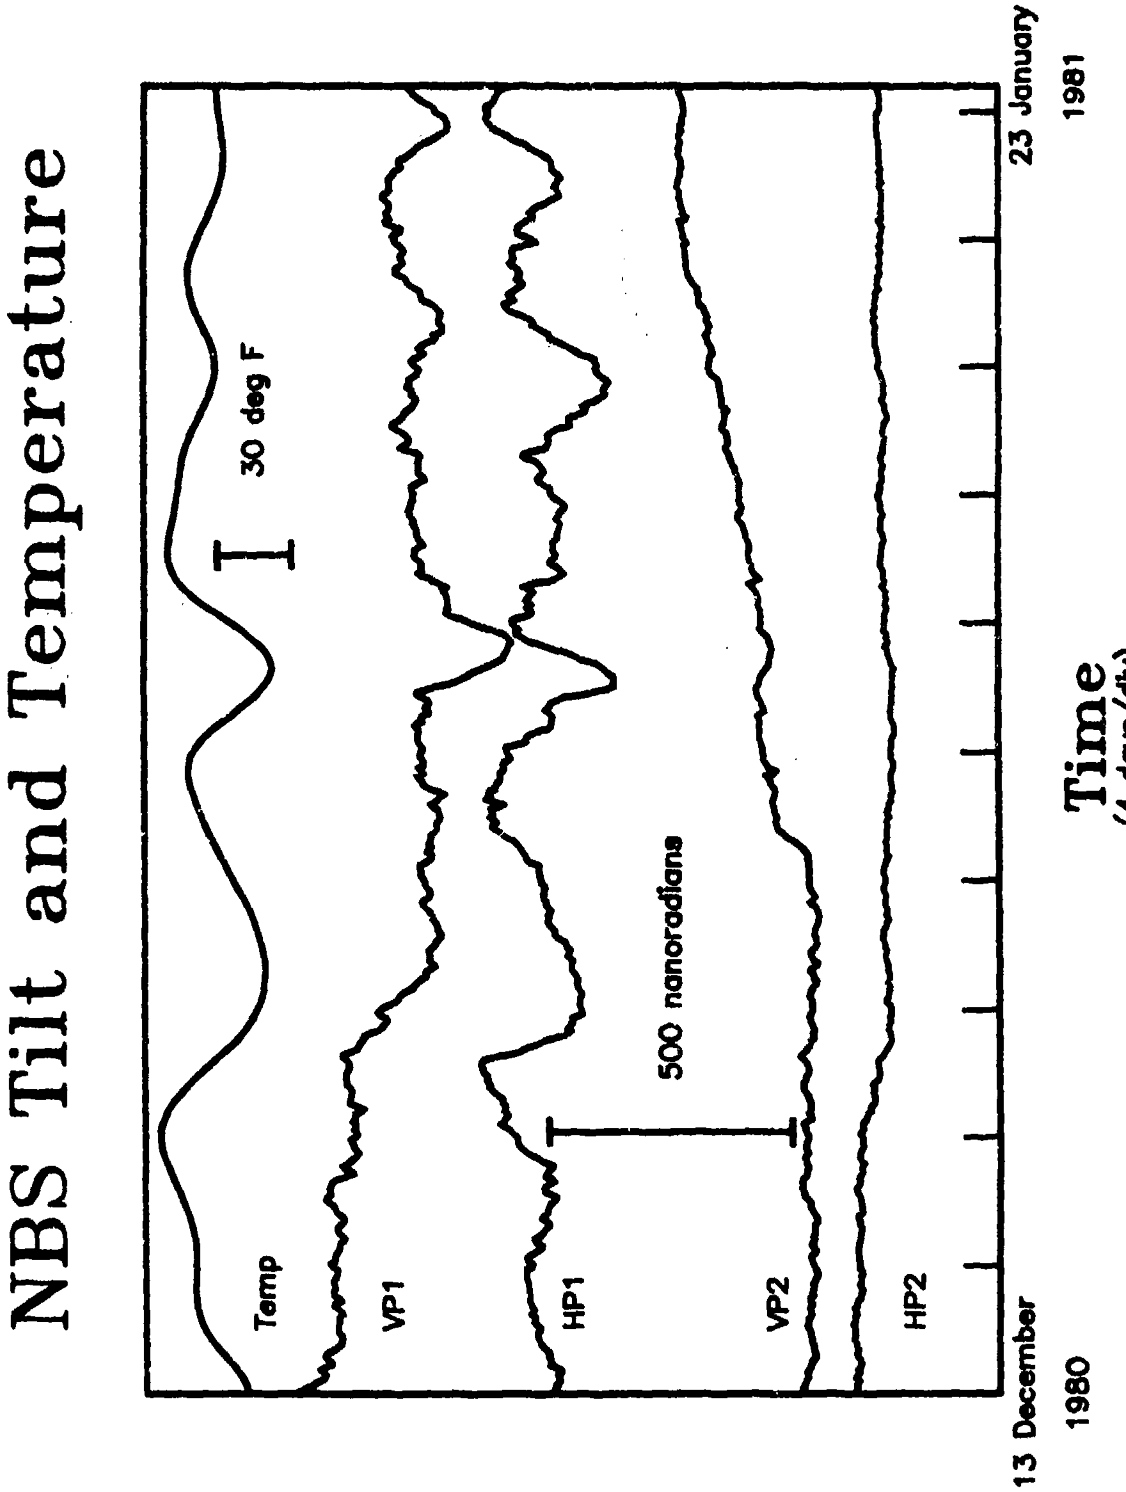

Fig. 7. Tidal residuals from four tiltmeters at the NBS site. The traces marked VP1 and VP2 are data from two vertical pendulums in a hole 16 meters deep. The traces marked HP1 and HP2 are data from two horizontal pendulums in a hole 33 meters deep. The top trace shows lowpassed surface temperature.

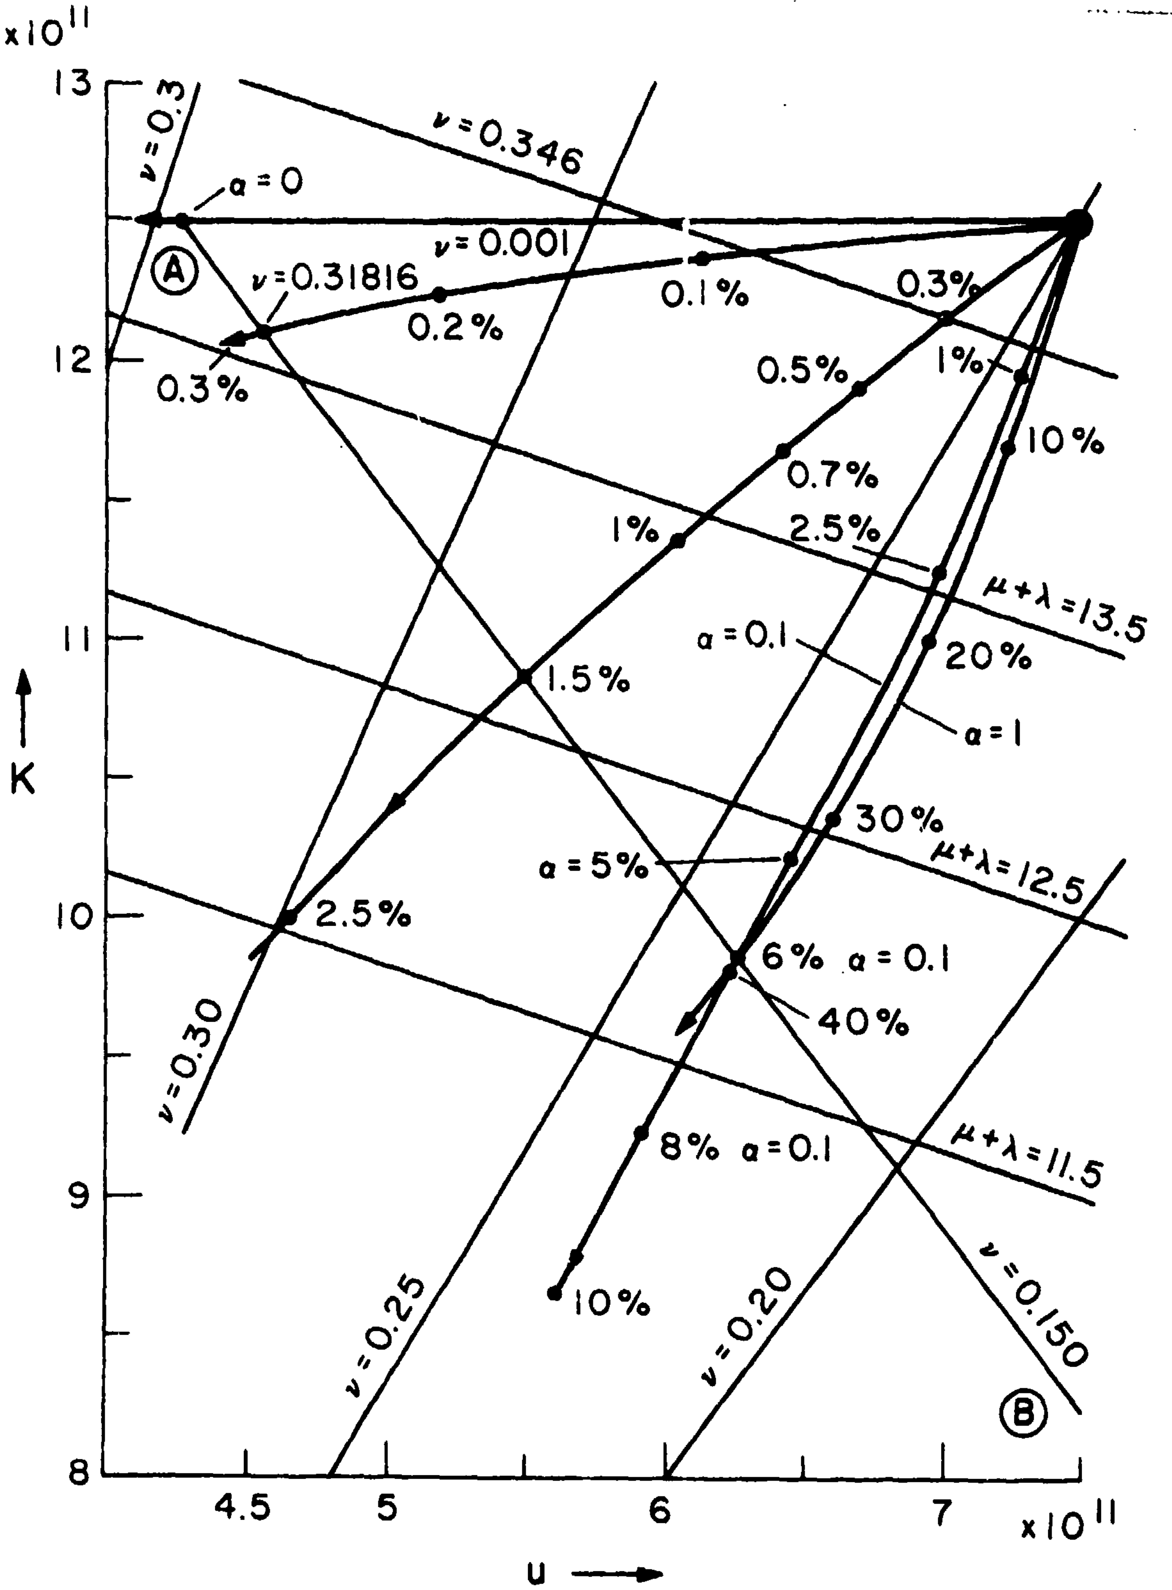

Fig. 8. K-U plot showing normal mantle material (Vp=8.13 km/s, density=3.4 gms/cc) at the large dot in the right hand corner. The locus AB corresponds to a 10% decrease in Vp. Trajectories are marked with the thickness/diameter ratio of the melt inclusions determined using Walsh's (1969) theory of partial melting.

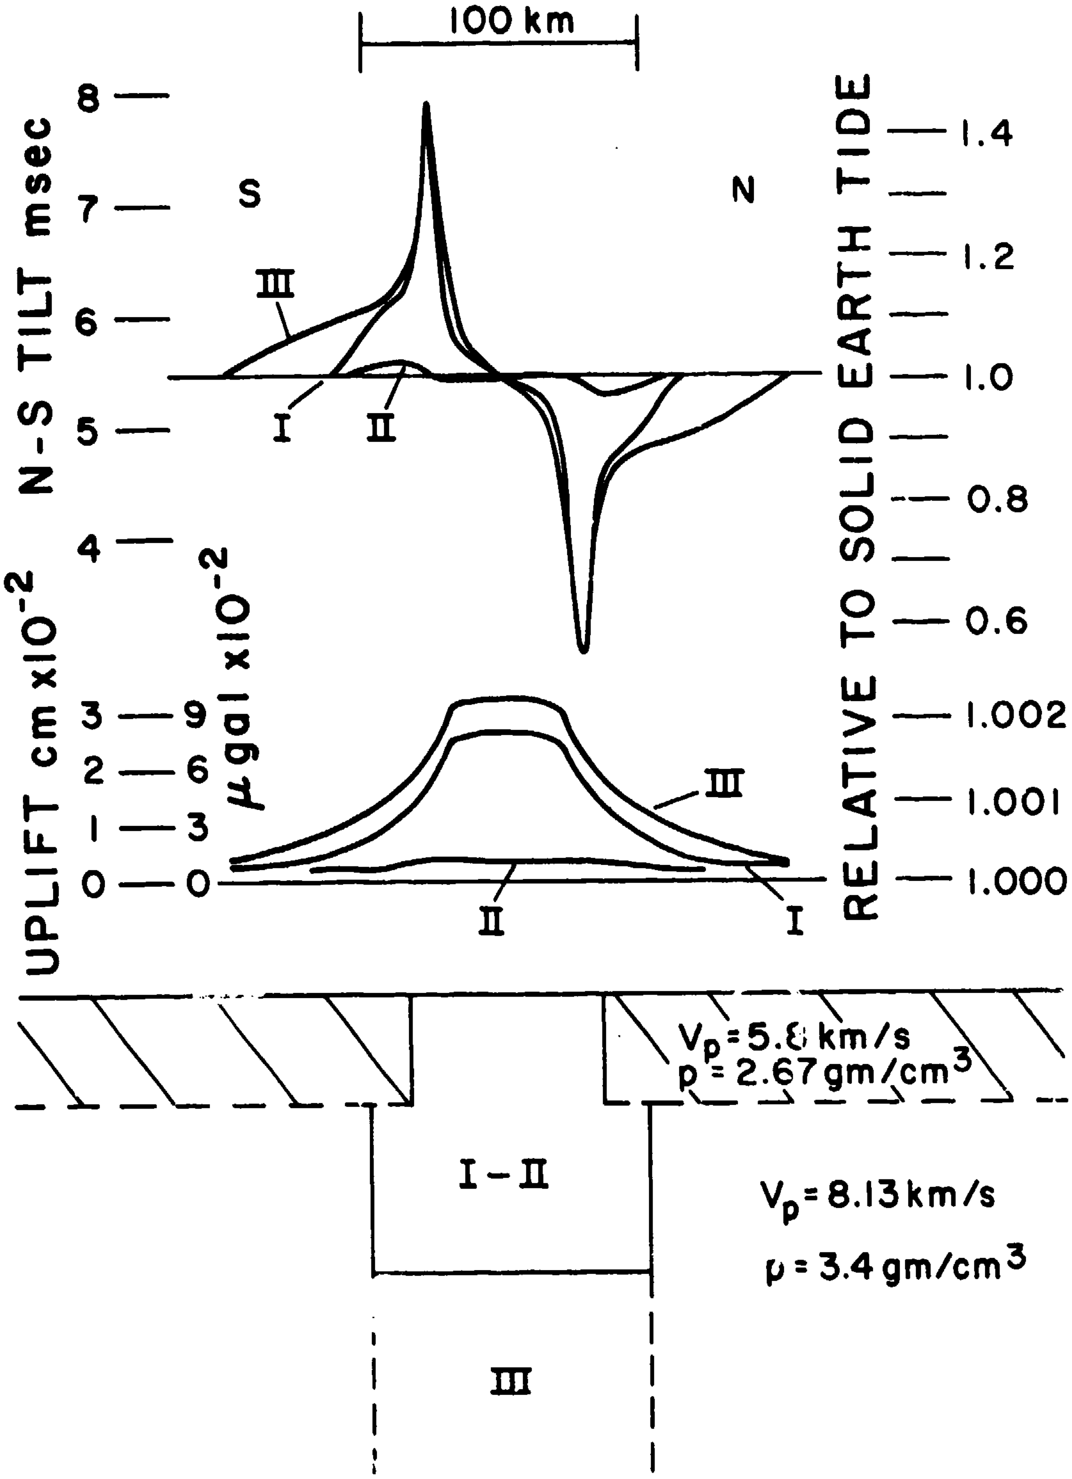

Fig. 9. Modifications in tilt, gravity and strain tides due to the partially molten zone beneath Yellowstone National Park assuming a 10% reduction in Vp. Models (I) and (II) are for the body shown extending to 100 km. In (I) the material properties correspond to flat inclusions. In (II) the properties correspond to round inclusions. Model (III) has the same material properties as (I) but extends to 200 km depth (taken from Harrison and Flach, 1976).

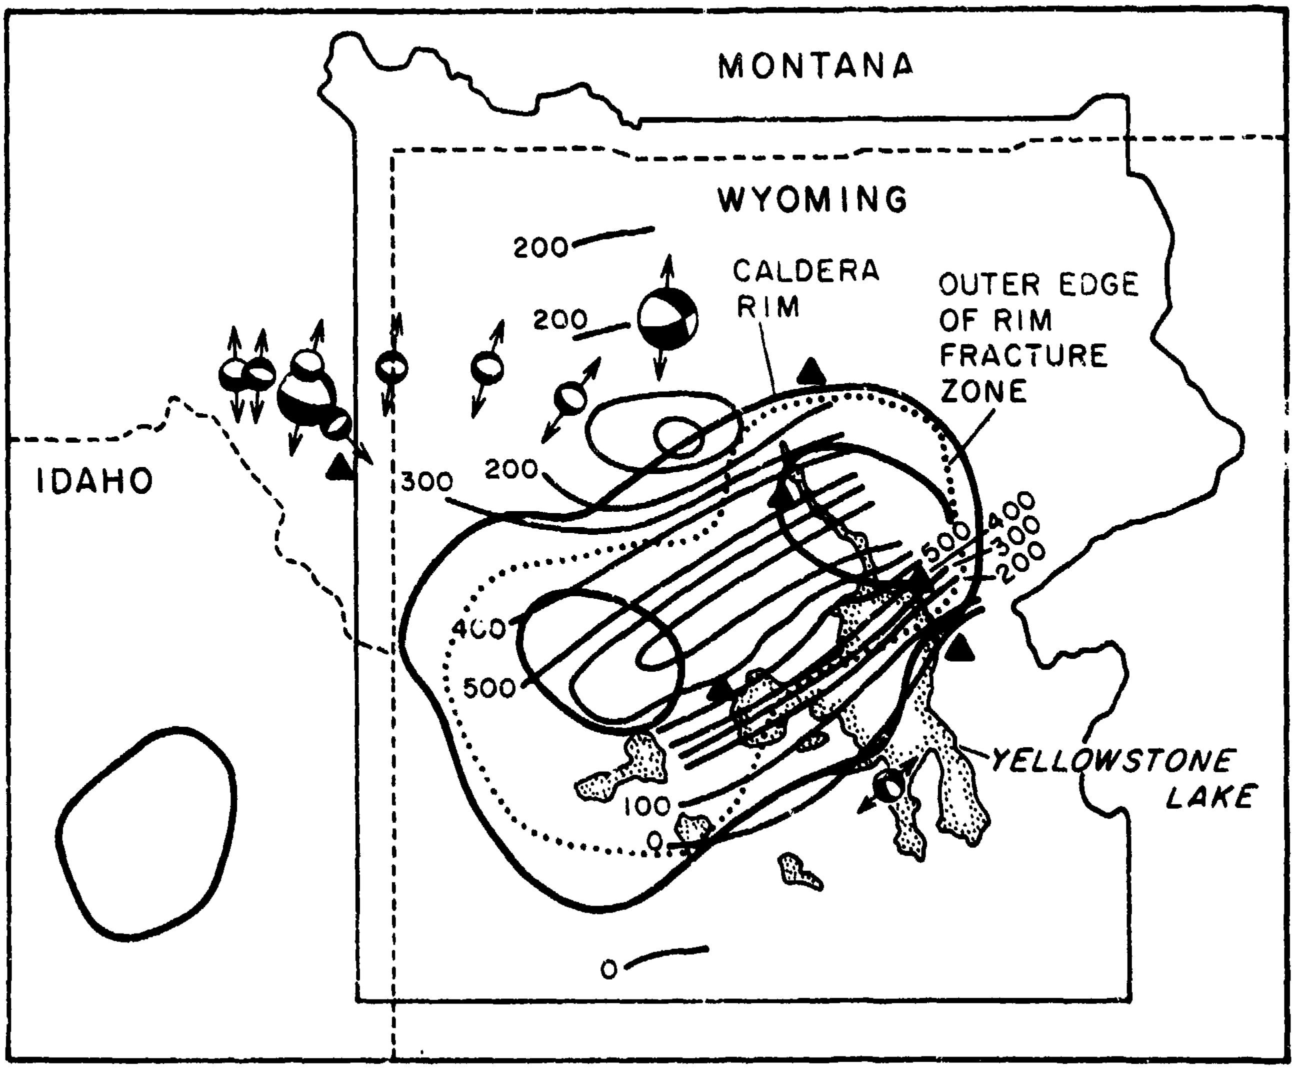

Fig. 10. Map of Yellowstone National Park showing uplift contours. Our proposed tiltmeter sites are shown by triangles.